BANK OF ENGLAND

Speech

The impact of machine learning and AI on the UK economy – conference overview

Comments delivered by

David Bholat

Senior Manager, Advanced Analytics

Conference on the impact of machine learning and AI on the UK economy

Bank of England

25 Mar 2020

These comments reflect my views and may not reflect those of the Bank of England. Thanks to my co-organisers: the Brevan Howard Centre at Imperial College London, particularly Franklin Allen; and the Centre for Economic Policy Research, especially Tessa Ogden, co-organiser Antonio Fatas, and Mandy Chan. I would like to also thank those who supported the organisation of the event: Paul Robinson; Tom Mutton; Misa Tanaka, Michael Kumhof; Neeltje Van Horen; Gosia Makarczuk; Verina Oxley; Candy Stephens; Juan Mateos-Garcia; Pete Thomas; Tangy Morgan; Andreas Viljoen; Tom Smith; and, especially, Michelle Scott.

I. INTRODUCTION

Cybersecurity. Climate change. And now coronavirus.

Some of the most critical issues impacting humans today concern our interdependence with the non-human environment.[1]

This conference addresses one such issue: the potential impact machines, imbued with artificial intelligence, may have on our economy, and the unique opportunities and challenges this presents to economic policy institutions like the Bank of England.

We had originally planned a conference hosted physically at the Bank. However, due to the Covid-19 pandemic, we are recording a series of video presentations instead. By doing so, we hope to make virtue out of necessity. Since the conference presentations are now video recorded, we are able to share their content with many more people than we could have physically accommodated at the Bank. And it is in fact fitting that a conference about the digital economy is delivered by digital means.

In these opening remarks, I would like to set the scene for the rest of the presentations by providing a high-level overview of the key issues around which they are structured.

II. THE HISTORY AND FUTURE OF ARTIFICIAL INTELLIGENCE

To begin, I'm delighted that Stuart Russell, Professor of Computer Science at UC Berkeley, is presenting the conference's keynote, providing an overview of the history and future of Artificial Intelligence. Artificial Intelligence or AI, as Stuart defines it in his leading textbook, refers to machines that receive perceptual input from the environment and perform actions in response.[2] Under this broad umbrella, three more specific but by no means mutually exclusive meanings are often intended.[3] The first is AI as short-hand for automation, including robotics. So understood, AI has been with us for a while. By contrast, a second meaning ascribed to AI refers to a possibility that does not yet exist. This is the stuff of sci-fi, what Nick Bostrom has termed 'superintelligence', when machine intelligence exceeds human intelligence in general, though this is already the case in certain specific tasks.[4] It can also refer to a speculative future Ray Kurzweil has called the 'singularity', a tipping point reached in the history of humanity through our steady accretion of prosthetics– from wearable smart glasses to the potential implanting of nanotechnologies– that so blurs the line between man and machine that we become a qualitatively different species from our ancestors.[5] Finally, more prosaically but more commonly, AI refers to machine learning, a set of statistical algorithms applied to data for making predictions or finding useful new patterns.

Machine learning emerged as an area of active research in computer science during the 1950s, for example, through the pioneering work of Alan Turing, who now features on the Bank of England's £50

[1] Smart, A. and Smart, J. 2017. *Posthumanism*. Toronto: University of Toronto Press.
[2] Russell, S. and Norvig, P. 2016. *Artificial Intelligence: a modern approach*. Harlow: Pearson.
[3] Agrawal, A., Gans, J. and Goldfarb, A. 2019. Introduction in *The economics of Artificial Intelligence: an agenda*, Agrawal, A., Gans, J. and Goldfarb, A., editors. Chicago: University of Chicago Press.
[4] Bostrom, N. 2014. *Superintelligence: Paths, dangers, strategies*. Oxford: Oxford University Press.
[5] Kurzweil, R. 2005. *The singularity is near: when humans transcend biology*. New York: Viking. See also Harari, Y. *Sapiens: A brief history of humankind*. London: Vintage.

note. However, only in the past decade or so, has machine learning reached a critical mass in the public conscious. Three key factors have been responsible. First, there has been an explosion in the volume, velocity and variety of data generated by digitalisation since the advent of the Internet. Data is the oil that powers AI—the more data machine learning algorithms are trained on, the more accurate their predictions typically become. Hence they are often associated with Big Data. Second, there has been a concomitant improvement in the technology to store these data and to analyse them efficiently. On-demand cloud computing, coupled with increased computational power owing to improvements in memory and processing speeds, have reduced the cost of data archiving.[6] At the same time, free, open-source software like R and Python with machine learning libraries have lowered barriers to analytical entry and popularised machine learning. Third, there have been improvements in the algorithms themselves. This has been most notable in the areas of deep learning and reinforcement learning. According to the World Intellectual Property Organization, while AI-related patents have been filed since the 1950s, over half have been published since 2013, with deep learning filings growing the fastest.[7]

III. THE 'FOURTH INDUSTRIAL REVOLUTION:' FACT OR FICTION?

Invention is one thing. Application is another. How significantly is machine learning and AI transforming the public and private sectors? This question is at the heart of a slate of presentations organised under the title 'The Fourth Industrial Revolution: Fact or Fiction?' Here, there are conflicting accounts. For some, machine learning and AI are general purpose technologies, that is, technologies like electricity and the Internet, which revolutionised industry across a range of sectors. What makes AI a general purpose technology is its ability to catalyse further innovation by more efficiently searching high-dimensional spaces to identify useful patterns and combinations.[8] Recently, for example, scientists at MIT discovered a new antibiotic by computationally combining chemical compounds at a scale unfeasible by physical experimentation in a laboratory.[9]

However, while some commentators claim we are in the midst of a Fourth Industrial Revolution, others, observing slow growth and low productivity gains for over a decade in the UK and other advanced economies, contrastingly claim that ours is an era of 'secular stagnation.'[10] Who's right?

To help us weigh up the balance of evidence for and against these opposing viewpoints, and illuminate possibilities that lie between them, we have a multidisciplinary group of presenters comprised of Manuela Veloso, a computer scientist and head of AI research at JP Morgan; Guy Michaels, an economist at the LSE; and Nick Craft, a historian with particular expertise in British economic history.

[6] Financial Stability Board. 2017. Artificial Intelligence and machine learning in financial services.
[7] World Intellectual Property Organization. 2019. Artificial Intelligence.
[8] Agrawal, A., McHale, J. and Oettle, A. 2019. Finding needles in haystacks: artificial Intelligence and recombinant growth in *The economics of Artificial Intelligence: an agenda*, Agrawal, A., Gans, J., and Goldfarb, A., editors. Chicago: University of Chicago Press.
[9] Sample, I. 20 February 2020. Powerful antibiotics discovered using machine learning for first time. *The Guardian*.
[10] Summers, L. 2014. Reflections on the new 'Secular Stagnation' hypothesis.

IV. THE IMPACT OF MACHINE LEARNING AND AI ON UK FINANCIAL SERVICES

One sector where machine learning and AI is making rapid inroads is in UK financial services. This is the subject of another group of presentations. To the point, a recent Bank of England and FCA survey of UK financial firms, which the Bank's Head of FinTech, Ashley Young, discusses in more detail in her presentation, revealed that the majority of respondents report they are already using machine learning. Applications run the gamut, from how firms model credit risk, to how they interact with customers through chatbots, as well as to how they detect fraud and execute trading opportunities. There is particular interest in how machine learning and allied data science techniques can automate regulatory reporting and compliance, an area dubbed 'RegTech.'

RegTech holds the promise of delivering efficiencies in a sector where some studies have suggested unit costs have remained stable for very long periods of time.[11] Yet these efficiency gains may not translate immediately into increased profitability. On the contrary, the cost curve for firms making substantial investments in machine learning and AI capabilities, and the complementary business processes to support them, could be upward-sloping, at least in the short-run. And if these strategic change projects fail, this amplifies the operational risks financial firms run.[12] We are therefore fortunate to have Louise Herring from McKinsey, and Ulku Rowe from Google, share insights based on their experience working with and for companies where AI has been successfully embedded, offering best practice tips and tricks on how financial institutions can mitigate downside risks and maximise their return on investment.

I am also delighted that Bonnie Buchanan, Head of the Department of Finance and Accounting at Surrey University, is a presenter. Bonnie authored a report for the Alan Turing Institute on AI in finance last year.[13] Her report is a salient reminder that while AI and ML are virtual technologies, they take shape in particular places, embodied in clusters like Silicon Valley and Shenzhen. Notably, while the UK invests in AI more than most advanced economies, the United States and China invest 50 times and eight times more, respectively.[14] Which countries win the race to become "AI superpowers"[15] may have potential long-term implications on the dynamics of international relations and on patterns of global trade.

V. ETHICAL AND CONSUMER CONDUCT ISSUES RAISED BY AI

What would have been Day 1 of our conference concludes with a series of presentations reflecting on some of the ethical and consumer conduct issues raised by artificial intelligence. This topic covers a potentially broad range of possible issues. These include the threat AI-generated deep fakes could pose to the quality of information consumers receive; the appropriate risk controls governing the development and deployment of algorithms in firms; how to balance the rights of individuals to privacy against

[11] Philippon, T. 2019. On fintech and financial inclusion.
[12] Bank for International Settlements. 2018. Implications of fintech developments for banks and bank supervisors.
[13] Buchanan, B. 2019. Artificial Intelligence in finance.
[14] McKinsey Global Institute. 2019. Artificial Intelligence in the United Kingdom: prospects and challenges.
[15] Lee, K. 2018. *AI superpowers: China, Silicon Valley and the new world order*. New York: Houghton Mifflin.

commercial interests in pricing insurance more precisely through sensors, GPS and other sources of data[16]; and remedying potential biases in data on which models are trained, when these data either misrepresent, or miss out entirely, particular demographic groups.

In the UK, financial conduct regulation, including the prevention of market abuse and ensuring consumers get a fair deal from financial firms, is the responsibility of the Financial Conduct Authority, so I am delighted that Karen Croxson, Head of Research and Deputy Chief Economist at the FCA, is a presenter. While not strictly within the PRA's regulatory remit, unethical behaviour and misconduct at firms are often leading indicators of risks to firms' safety and soundness, which is an objective.[17] As firms adopt AI, the link between conduct and prudential issues may become even closer. For example, it is not difficult to imagine possible incidents involving the mishandling of data, or reports of biased algorithms, leading to depositors and investors losing confidence in a firm.

The widespread adoption of machine learning algorithms across the financial sector could also have important distributional consequences for consumers, for example, in terms of credit allocation. Hopefully, the combination of machine learning and new sources of data means those who may have been financially excluded in the past now receive credit in the quantity and at the price that reflects their true underlying credit risk. Nevertheless, other individuals may find that machine learning models make them less well off than if traditional modelling approaches had been employed. Research presented by one of our speakers, Tarun Ramadorai, Professor of Financial Economics at Imperial College London, indeed suggests that machine learning algorithms can have these Janus-faced effects, notably among historically disadvantaged groups.

One implication of Tarun's findings is that it may be wise for other goals besides predictive accuracy, such as fairness, to be explicitly incorporated into the objective function of machine learning models. A major challenge in doing so is how to operationalise a concept like fairness. In his recently published book, *The Ethical Algorithm*, one of our presenters, Michael Kearns, offers guidance on how this might be done.[18] Yet, as Michael points out, even if we can plot the space of optimal trade-offs between fairness and accuracy, where on the trade-off curve we want to be is ultimately a decision for human, not artificial, intelligence.

VI. MACHINE LEARNING, AI AND FINANCIAL STABILITY

Day 2 of the conference would have begun with a panel on machine learning, AI and financial stability. As elsewhere, here there is a healthy debate to be had about the direction and magnitude of the effect. For example, Joel Suss, a researcher at the Bank and a presenter, has found that machine learning models more accurately predict UK bank distress than standard models.[19] At the same time, Jon Danielsson, Director of the Systemic Risk Centre at LSE, and another of the presenters, has expressed

[16] Patel, K. and Lincoln, M. 2019. It's not magic: weighing the risks of AI in financial services.
[17] Bailey, A. 2016. Culture in financial services– a regulator's perspective.
[18] Kearns, M. and Roth, A. 2020. *The ethical algorithm: the science of socially aware algorithm design*. Oxford: Oxford University Press.
[19] Suss, J. and Treitel, H. 2019. Predicting bank distress in the UK with machine learning.

concerns that if firms adopt the same best performing models trained on similar data, then this introduces a new form of conformity and set of positive cross-correlations in the financial system.[20] If the system experiences a negative shock, then firms, taking their steer from these models, could behave the same way through sell-offs that coalesce into systemic risk.

The presentations of Joel and Jon are complemented by those from Sheri Markose, Professor of Economics at Essex University, and Andy Haldane, the Bank's Chief Economist. Each draws on an array of natural and social science fields beyond economics to highlight the comparative strengths and weaknesses of human and artificial intelligences. A key take-away from their presentations for me is that we should remain cautious about how much machine learning models trained on past data can tell us about the future path of financial stability. The areas where AI and machine learning have been most fruitfully applied, such as facial recognition or winning the game Go, involve prediction targets that are static, and where the rules are fixed and finite. This is not the case in financial markets, which are dynamic, non-stationary domains. Furthermore, there is an iterative and reflexive form of co-determination between models and the economy that doesn't exist in non-social science applications; as per the Lucas critique, actions based on model predictions can change previously observed statistical patterns.

VII. NEW COMPETITIVE DYNAMICS CREATED BY MACHINE LEARNING AND AI

Machine learning and AI also have the potential to change the market structure of financial services, potentially making them either more, or less, competitive. We have assembled a group of experts to discuss these issues: Giacomo Calzolari, Professor of Economics at the European University Institute; Paul Grout, Professor of Political Economy at Bristol University and Senior Adviser for Competition at the Bank of England; and Kate Collyer, Chief Economist at the FCA, who, unfortunately, but quite understandably, hasn't been able to participate due to current exigencies.

There are certainly reasons to think AI and machine learning might facilitate effective competition in financial services, which is a secondary objective of the Bank's PRA. As mentioned earlier, machine learning software and the compute power for Big Data analytics is either free or cheaply available, suggesting that financial markets could become more contestable by new entrants. Unlike incumbents, new entrants may be able to move more nimbly because they are unburdened by legacy systems and existing IT infrastructural dependencies.

However, a striking feature of the digital economy is that while machine learning and AI algorithms are mostly free and open-source, the financial data on which they can be trained is typically closed and proprietary.[21] Thus existing incumbents, with longer loan books and possibly deeper pockets to invest in data science talent and supporting infrastructure, could see their market power enhanced; by more

[20] Danielsson, J., Macrae, R., and Utheman, A. 2019. Artificial Intelligence and systemic risk.
[21] Cockburn, I., Henderson, R., and Stern, S. 2019. The impact of artificial intelligence on innovation in *The economics of Artificial Intelligence: an agenda*, Agrawal, A., Gans, J. and Goldfarb, A., editors. Chicago: University of Chicago Press.

efficiently exploiting their Big Data assets with techniques like machine learning, they might be able to deliver improvements in the quality of the services they offer consumers, winning them even more, making the financial services sector more concentrated and less competitive.[22]

In policy circles, particular attention has attached to whether Big Tech firms such as Amazon, Facebook or Google might expand their financial services offerings. For example, both the BIS and the Financial Stability Board have analysed this issue.[23] While increasing competition in the near-term, Big Tech firms' greater entrance into financial services could lead to greater concentration over longer time horizons.

Besides Big Tech, other competition issues loom. For example, and anticipating Giacomo's presentation, there is some computationally simulated experimental evidence to suggest that certain classes of reinforcement algorithms might be able to learn to tacitly collude with each other to charge supra-competitive prices without being explicitly programmed to do so.[24] How to ensure financial markets are fair and effective when agents in those markets are algorithmic actors could become a key policy challenge.

VIII. LABOUR MARKET AND REAL ECONOMY IMPLICATIONS

We then turn to thinking about the impact of machine learning and AI on the real economy. Like many others, this set of contributions centres on a paradox. On the one hand, there are many observers who forecast greater economic growth driven by AI. According to McKinsey, AI could deliver additional economic output of around $13 trillion globally by 2030.[25] In the UK alone, it has been estimated that AI could add an additional £630 billion to the economy over the next fifteen years.[26] On the other hand, as Erik Brynjolfsson and his co-authors have quipped, with a nod to Nobel Prize-winning economist Robert Solow, we are seeing the transformative impact of AI and machine learning 'everywhere but in the productivity statistics.'[27]

One possibility is that we are measuring the impact of new techniques and technologies imperfectly. As I mentioned earlier, core machine learning tools, like R and Python, are free, so their economic value may not be captured adequately by GDP statistics premised on goods and services with market prices. More generally, home-grown intangible investments such as the unique data a firm possesses, and the machine learning models firms code to exploit these data, don't appear on firms' balance sheets. This makes valuation of individual firms and investments in the economy as a whole more difficult in the era

[22] Farboodi, M., Mihet, R., Philippon, T., and Veldkamp, L. 2019. Big Data and firm dynamics.
[23] Carstens, A. 2018. Big tech in finance and new challenges for public policy. Financial Stability Board. 2019. BigTech in finance: market developments and potential financial stability implications.
[24] Calvano, E., Calzolari, G., Denicolo, V., and Pastorello, S. 2019. Artificial Intelligence, algorithmic pricing and collusion.
[25] McKinsey Global Institute. 2018. Notes from the AI frontier: modelling the impact of AI on the world economy.
[26] Hall, W. and Presenti, J. 2017. Growing the artificial intelligence industry in the UK.
[27] Brynjolfsson, E., Rock, D., and Syverson, C. 2018. The productivity J-curve: how intangible capital complements general purpose technologies.

of AI. In a recent book and in his presentation, Jonathan Haskel, a member of the Bank's Monetary Policy Committee, explores mismeasurement as a potential explanation for the productivity paradox.[28]

Another contributory factor could be that most workers have not yet acquired the skills, knowledge and experiences to efficiently exploit new techniques and technology. Indeed, many discussions on the economic impact of machine learning and AI focus on the labour market. Both Alan Manning, Professor of Economics at the LSE, and Eduardo Rodriguez Montemayor, an economist at PWC, deliver presentations illuminating this issue.

The impact of machine learning and AI on the labour market could be multifaceted, effecting the quantity of jobs and their nature, as well as their remuneration. Daniel Susskind of Oxford University, who provides a response to Jonathan, Alan, and Eduardo's presentations, has written about the possibility of 'a world without work'— the title of his new book.[29] The basic idea is that large-scale implementation of AI and automation could lead to higher levels of structural unemployment. Even if that does not happen, some existing occupations are likely to disappear, just as entirely new occupations are created. And those existing occupations that remain may be transformed. Famously, the invention of the ATM counterintuitively increased the employment of bank tellers, but changed bank tellers' roles from cashiers to sales associates.[30]

In this context, it is also worth recalling that machine learning and AI will not only impact the economy, but that economic fundamentals of supply and demand, relative scarcity and comparative advantage will influence the pace and places where machine learning and AI develop. In particular, as during the First Industrial Revolution, the extent to which there is a substitution toward capital and away from labour will be influenced by the relative quantities and prices of each.[31] And among workers, the benefits of the Fourth Industrial Revolution may be bestowed unevenly, with some commentators suggesting it could exacerbate income inequalities, while others suggest any relative nominal income losses might be offset by real gains for all workers in their guise as consumers of higher quality and lower priced goods and services.[32]

IX. APPLYING MACHINE LEARNING AND AI IN CENTRAL BANKING AND REGULATION

Our final session considers how machine learning and AI can impact economic policymaking, with contributions from Dave Ramsden, Deputy Governor at the Bank of England; James Proudman, Senior Adviser in the PRA; and Gareth Ramsay, Executive Director of the Data Analytics Transformation Directorate.

[28] Haskel, J. and Westlake, S. 2018. *Capitalism without capital: the rise of the intangible economy*. Princeton: Princeton University Press.
[29] Susskind, D. 2020. *A world without work: technology, automation, and how we should respond*. London: Allen Lane.
[30] Bessen, J. 2016. How computer automation affects occupations: technology, jobs and skills.
[31] Allen, R. 2009. *The British Industrial Revolution in global perspective*. Cambridge: Cambridge University Press.
[32] Cowen, T. 2019. Neglected open questions in the economics of artificial intelligence in *The economics of Artificial Intelligence: an agenda*. Agrawal, A., Gans, J., and Goldfarb, A., editors. Chicago: University of Chicago Press.

Machine learning, especially supervised algorithms, are powerful additions to the central banker's toolkit. They often outperform traditional statistical approaches in prediction tasks because they can detect non-linear patterns in data, and handle datasets with a large number of predictors. Yet unlike basic statistical models, their internal complexity means it is not often easy to ascribe clear causal meaning to specific variables.[33] Given the importance which central banks like the Bank of England attach to public transparency and accountability, the opacity of machine learning models poses a challenge.

A promising stream of new Bank research is addressing this challenge by making machine learning models more explainable.[34] And machine learning models are being increasingly applied for various analytical purposes at the Bank, from modelling the probability of financial crises[35] to understanding how supervisors write to firms.[36]

A particular area of investment is in using machine learning, and, more broadly, data science techniques, to aid prudential supervision– so-called SupTech. This is work led by our PRA Data Innovation team. For example, James Proudman has spoken of an evolving form of 'cyborg supervision' involving humans and machines "working ever more closely together and leveraging their comparative strengths."[37] The crucial point James stresses is that for the foreseeable future, the PRA's supervisory regime will remain firmly anchored in human judgment, but that over time this judgment may be exercised in new domains such as how models are built. At least for now, machine learning models very much require a human-in-the-loop to make judgments about which algorithms to use; the prediction target; what predictors to include; and how to measure model performance.[38] And while these models may inform supervisor's forward-looking judgement, key decisions are unlikely to be automated, but rather continue to rest with supervisors, who can consider the outputs from machine learning models alongside other information that is not easy to quantify or incorporate into them. For central banks and regulators, it is therefore important that machine learning and AI remains 'human compatible,' to echo the title of Stuart Russell's recent book.[39]

[33] Mullainathan, S. and Spiess, J. 2017. Machine learning: an applied econometric approach.
[34] Joseph, A. 2019. Shapley regressions: a framework for statistical inference on machine learning models. Bracke, P., Datta, A., Jung, C., and Sen, S. 2019. Machine learning explainability in finance: an application to default risk analysis.
[35] Bluwstein, K., Buckmann, M., Joseph, A., Kang, M., Kapadia, S., and Simsek, O. 2020. Credit growth, the yield curve and financial crisis prediction: evidence from a machine learning approach.
[36] Bholat, D., Brookes, J., Cai, C., Grundy, K., and Lund, J. 2017. Sending firm messages: text mining letters from PRA supervisors to banks and building societies they regulate.
[37] Proudman, J. 2018. Cyborg supervision– the application of advanced analytics in prudential supervision.
[38] Hunt, S. 2017. From maps to apps: the power of machine learning and artificial intelligence for regulators
[39] Russell, S. 2019. *Human compatible: AI and the problem of control*. London: Allen Lane.