

Machines, jobs and equality

**Technological change and
labour markets in Europe**

Edited by Andreas Bergström and Karl Wennberg





Andreas Bergström is Deputy Director at the green and liberal think tank Fores (Forum for Reforms, Entrepreneurship and Sustainability) and head of Fores' programme for economic reforms and entrepreneurship. He formerly worked as Political Adviser for several liberal ministers, and now leads research projects on automatisisation, entrepreneurship and employment.



Anna Breman is Group Chief Economist at Swedbank since 2015, while also serving as independent expert in the government analysis group "The future of Work," holding a Ph.D in Economics from Stockholm School of Economics. She previously worked as deputy director at the Swedish Ministry of Finance, and as a researcher in the US.



Fernando del Río is Assistant Professor of Economics at the Universidade de Santiago de Compostela, holding a Ph.D. in Economics from Universidad Carlos III de Madrid. His areas of research are economic growth and labour economics.



Eduardo L. Giménez is Assistant Professor of Economics at the Universidade de Vigo, with a Ph.D. in Economics from Universidad Carlos III de Madrid. His areas of research are, among others, general equilibrium theory and its applications to real business cycles, financial economics and monetary theory.



Georg Graetz is an Assistant Professor in the Department of Economics at Uppsala University, research affiliate at the Institute for the Study of Labor (IZA), and associate at the Centre for Economic Performance in London. His research focuses on technological change, wage inequality, and the economics of education.



Claire Ingram is a PhD Candidate at the Stockholm School of Economics in Stockholm, Sweden. She conducts research into the impact of increased reliance on digital artefacts, especially when it comes to entrepreneurship and the world of work.



Darja Isaksson is an entrepreneur and advisor within consumer experience design, innovation strategy and data analytics. She was awarded 'Most influential opinion maker of the year' in 2016 by Swedish business magazine Veckans Affärer, and is currently a member of the Swedish Prime Minister's innovation council.

Machines, jobs and equality

Andreas Bergström and Karl Wennberg (editors)

First edition, first print

Printed by Fridholm & Partners, Hindås, Sweden, 2016

ISBN: 978-91-87379-36-9

Published by the European Liberal Forum asbl with the support of the Bertil Ohlin Institute and Fores. Funded by the European Parliament. The European Parliament is not responsible for the content of the publication. The views expressed are those of the authors alone. They do not necessarily reflect the views of the European Liberal Forum asbl.

Design by Studio Annalog

© 2016 The European Liberal Forum (ELF). This publication can be downloaded for free on www.liberalforum.eu or www.fores.se. We use Creative Commons, meaning that it is allowed to copy and distribute the content for a non-profit purpose if the author and the European Liberal Forum are mentioned as copyrights owners. (Read more about creative commons here <http://creativecommons.org/licenses/by-nc-nd/4.0/>)

The European Liberal Forum (ELF) is the foundation of the European Liberal Democrats, the ALDE Party. A core aspect of our work consists in issuing publications on Liberalism and European public policy issues. We also provide a space for the discussion of European politics, and offer training for liberal-minded citizens. Our aim is to promote active citizenship in all of this. Our foundation is made up of a number of European think tanks, political foundations and institutes. The diversity of our membership provides us with a wealth of knowledge and is a constant source of innovation. In turn, we provide our members with the opportunity to cooperate on European projects under the ELF umbrella. We work throughout Europe as well as in the EU Neighbourhood countries. The youthful and dynamic nature of ELF allows us to be at the forefront in promoting active citizenship, getting the citizen involved with European issues and building an open, Liberal Europe.

Bertil Ohlininstitutet is a Swedish think tank, founded in 1993 with the purpose to initiate research and debate in critical areas of public policy in the tradition of liberal thinking. The institute is named after Bertil Ohlin, who was leader of the Liberal party of Sweden (Folkpartiet) between 1944 and 1967. By building networks of scholars, participants in the public debate and persons from private and public working life the institute – as a think tank – contributes to broadening the basis for liberal opinion formation and renewal of liberal thinking.

Fores is a green and liberal think tank. We are a foundation and non-profit NGO and we want to renew the debate in Sweden with a belief in entrepreneurship and creating opportunities for people to shape their own lives. Market-based solutions to climate change and other environmental challenges, the long-term benefits of migration and a welcoming society, the gains of increased levels of entrepreneurship, the need for a modernization of the welfare sector and the challenges of the rapidly changing digital society – these are some of the issues we focus on. We act as a link between curious citizens, opinion makers, entrepreneurs, policymakers and Z

Index

- 7 Introduction**
Andreas Bergström and Karl Wennberg
- 13 The maid, the clerk, the doctor & their computers**
Fernando del Rfo and Eduardo L. Giménez
- 51 The impact of technological change on the labour market**
Georg Graetz
- 65 Digitalisation and collective value creation**
Darja Isaksson and Karl Wennberg
- 87 Diginomics and the productivity puzzle**
Anna Breman
- 105 How may robots affect the labour market in the near future?**
Fredrik Löfgren
- 135 Building blockchains: In search of a distributed ledger ‘standard’?**
Claire Ingram, Jacob Lindberg and Robin Teigland
- 167 The intelligence explosion revisited**
Karim Jebari and Joakim Lundborg

Introduction

Introduction

The past few years have witnessed an intense debate regarding technological change and its effects on labour markets. Erik Brynjolfsson and Andrew McAfee's celebrated book "The Second Machine Age" and a host of related papers and books discuss how technological change related to computers and robots affects companies, labour markets, the distribution of wealth and entrepreneurship. While much of this work has been US-centric, European academics, policymakers and companies are increasingly discussing how the same issues affect European economies. Understanding these changes is a vital prerequisite to sound policy formation in the face of the societal challenges brought about by such rapid change, regardless of whether one adheres to the positive or negative scenarios outlined.

In an earlier book, "Inclusive Growth in Europe" (2014), we discussed the role of technological change in the rise of self-employment and 'precarious' work arrangements in Europe. In this book, we describe a number of important sources of technological change, such as information and computer technologies, accelerated robotisation, artificial intelligence, and the advent of new modes of payments, such as cryptocurrencies (Bitcoin). These are important and inter-related trends, fuelled by the rapid advancement in digital technologies and decreasing cost of computing power. The seven chapters in this volume

discuss the nature and significance of recent technological change, and the impact on European companies and labour markets. The first half of the book discusses the implications from a macro perspective, while the second half looks at the technologies.

Chapter one - *The maid, the clerk, the doctor and their computers* - by Fernando del Río and Eduardo Giménez, provides a review of recent research regarding technological change and labour markets, with a particular emphasis on Europe. Del Río and Giménez discuss research showing how information and communication technologies have enabled the automation and offshoring of many routine tasks. While this has contributed to increased earnings for high-skilled workers, it has also led to increasing job polarisation in the labour markets of most developed countries. This polarisation, the authors argue, poses new challenges to all democratic societies, as evidenced by segments of the population increasingly calling for populist political solutions in the face of economic decline. In their conclusions, the authors highlight income redistribution, improving educational systems and job training as ways of meeting these challenges.

The second chapter by Georg Graetz - *The impact of technological change on the labour market* - discusses the implications of recent technological change for economic development. Graetz - a leading labour market researcher - summarises research on how information technology has contributed to overall productivity growth and led to well-known changes in labour market demand for various skills and tasks. He also discusses what these developments could mean in the future. Specifically, he argues governments need to ensure the education system is responsive to the changing demands placed on workers, given the skills acquired in their youth may be outdated before they reach retirement age. Graetz shows that training for a particular occupation has become an uncertain investment, since automation could make this occupation more or less redundant. Governments could supply insurance for this type of risk by, for example, introducing re-training subsidies.

The third chapter by Darja Isaksson and Karl Wennberg - *Digitalisation and collective value creation* - describes how digitalisation and glo-

balisation interact as reinforcing trends. They discuss the impact of digitalisation for both industries and the public sector. In the private sector, it enhances competition by lowering barriers to funding, marketing, sales and distribution. However, Isaksson and Wennberg also warn that the large digital platforms providing data and standardising distribution mean one or a few of these can become dominant, as has already happened in the media and music industries. The authors discuss the dangers for competition and public welfare if standardisation leads to new monopolies, and the potential for digitalisation to leverage increased value in the production and distribution of collective goods, such as health care and energy.

The fourth chapter by Anna Breman - *Diginomics and the productivity puzzle* - discusses the conundrum that technological change (in the form of digitalisation) is not evidenced by enhanced productivity data. Breman outlines three potential explanations for this: that digital innovation is merely hype and not comparable to history's previous technological breakthroughs, that its effect is underestimated because productivity measures are not well adapted to assessing the impact of digital innovation, or that the productivity effects will come, but with a significant time lag. She presents arguments in favour of all three explanations, specifically for the lag effect, which she argues, provides the most plausible explanation for why productivity growth has tended to wax and wane throughout the last century. Breman concludes that regulatory changes are needed to further spur innovation and investment in digital technologies.

To understand the structural changes resulting from rapid technological change, one must also have some grasp of the technology involved. The second half of this book describes three areas we think are particularly important. The fifth chapter is written by Fredrik Löfgren, a prominent robot constructor and speaker. In *How may robots affect the labour market in the near future?* Löfgren outlines his experiences in the increasingly rapid development of robots, highlighting the existing potential for robots in manufacturing, arts, and interaction. Löfgren argues that the rate of development in robotisation is rapidly increasing and robots already exist with capabilities that would surprise most

people. He argues that in order to productively discuss the role of robots in today and tomorrow's society, it is necessary to accept the adverse as well as the positive effects robots will have.

The sixth chapter - *Building blockchains: In search of a distributed ledger 'standard'?* - by Claire Ingram, Jacob Lindberg, and Robin Teigland, depicts blockchain technology and its potential to fulfil an important role in electronic commerce by enabling 'the digitalisation of trust'. Ingram and co-authors describe the rapid evolution of blockchain technology as an emerging ecosystem containing different types of blockchains with the potential for competing content standards as well as competing sources of control. They argue that we know from experience that competition between standards in new technology requires careful handling by the authorities, but that self-regulation often occurs as 'standard wars' play out.

In the seventh and final chapter - *The intelligence explosion revisited* - Karim Jebari and Joakim Lundborg discuss the potential risks associated with artificial intelligence (AI). In their outline of current techno-philosophical discussions, Jebari and Lundborg highlight the distinction between two distinct AI-related risks - tool risk and agent risk - arguing that the former poses the greater risk. AI as autonomous agents are less a threat to society than the potential that AI will lead to a Hobbesian society where the introduction of too much complexity in global systems or the emergence of totalitarian surveillance states leads to the reduction of freedom for the many.

These seven chapters in this volume enable us to draw a number of conclusions. First, societal changes due to rapid technological change in the early 21st century are not as fast and frightening as sometimes described in the popular press. There is no good evidence that technological change so far has raised the unemployment rate. However, there is evidence to suggest that technological change has shifted the distribution of job types and productivity differences among occupations and sectors in the economy. These trends are likely to continue.

It is possible that in future it will be difficult for many individuals on the labour market to keep up with such rapid technological developments. High-quality education and continuous training of the work-

force will likely be a key policy tool to meet the challenges from the types of technological change we discuss. Leading theorists, such as Erik Brynjolfsson and Andrew McAfee, argue that schools need to focus more on enhancing students' creativity and willingness to learn new concepts and skills. Given the problems noted in primary and secondary school outcomes for several European nations, we cannot however dispense of the necessity that the future workforce also needs solid knowledge in verbal and writing skills, as well as mathematics. A broad general education is a necessary basis for workers' ability to adjust and benefit from subsequent education and training.

Related to the ever-increasing importance of public education keeping up with digitalisation are questions about redistribution. Economic inequality is a natural result of the type of rapid technological change described in this book. Questions about redistribution of wealth created by new innovations that destroy jobs will become more prevalent in the years to come. The research on job polarisation discussed in this volume's first two chapters clearly suggest we will see more of the kinds of jobs that are either highly paid or low-paid, and fewer 'good' jobs requiring intermediate education and training. While job polarisation is affecting the entire OECD community, stagnating median wages is more of a specific US phenomenon. To date, we are not seeing significant technology-driven unemployment in Europe. The debate around this kind of unemployment is largely influenced by US figures. In the EU, employment is quite high with the exception of a few countries still struggling with the aftermath of the great recession of 2007.

The discussion of digitalisation in industries, robots, cryptocurrencies and artificial intelligence we provide in this book highlight the importance of regulators legislating around such new technology. The future use of blockchain technology will depend on what is legally sanctioned and what is not. The same applies to new types of digital platforms selling goods and services in unforeseeable ways, sometimes loosely summarised as the sharing economy. Further, the discussion on artificial intelligence (AI) and its potential impact on society has barely begun. It is perfectly possible to imagine the accelerating development of AI leading to labour market changes, such that people without

Andreas Bergström and Karl Wennberg

specific skills and abilities difficult or expensive to replace with robots will not be able find paid employment. This may take 20 or 70 years. At any rate, robots are here to stay and discussions are needed to consider where we are heading as a society. Policy efforts regarding AI beyond those focusing on AI risks are however likely premature at this stage.

By discussing sources of recent technological change and how these affect European companies and labour markets, our goal with this book is to awaken interest and spur discussion among those interested in policy related to technology, competition and market effectiveness and labour markets. Providing a macro perspective first and subsequently a set of descriptions of how new technologies affect industries, financial transactions and labour markets, the book contains food for thought for politicians, managers, and academics in Europe and beyond.

We wish to thank ELF members, Mr. Hans Van Mierlo Stichting, Magma and the Novum Institute, for their invaluable help with the workshops that paved the way for this book. We also wish to thank all the authors, the anonymous referees, research editor Annalisa Tulinpano, Tove Mellgren who worked in the project, first as an employee of the Ohlin Institute and then at Fores, Jeroen Dobber at ELF, and Andreas Bergström's intern Gustav Juntti.

Andreas Bergström

Deputy Director, Fores

Karl Wennberg

Professor and Board Member, The Bertil Ohlin Institute

Fernando del Río is Assistant Professor of Economics at the *Universidade de Santiago de Compostela*, holding a Ph.D. in Economics from *Universidad Carlos III de Madrid*. His areas of research are economic growth and labour economics. **Eduardo L. Giménez** is Assistant Professor of Economics at the *Universidade de Vigo*, with a Ph.D. in Economics from *Universidad Carlos III de Madrid*. His areas of research are general equilibrium theory and its applications to real business cycles, financial economics and monetary theory.

The maid, the clerk, the doctor & their computers

The impact of information and communication technologies on jobs and wages

Contrasting fortunes

At the end of the eighties the British film director, Peter Greenaway, wrote and directed his most successful film entitled, *The Cook, the Thief, His Wife & Her Lover*. Greenaway narrates the intertwined stories of the four characters referred to in the title of his film. Along similar lines, this chapter also aims to tell the stories of four characters. Of those four, three of them are made of flesh and blood; the other is the outcome of human ingenuity. The fate of the first three characters is closely linked to the evolution of the latter. ‘The maid’ is the alter-ego of *low-*

skill, low-wage workers, ‘the clerk’ of middle-skill, middle-wage workers, and ‘the doctor’ of high-skill, high-wage workers. ‘The computer’ represents the role of the information and communication technologies (ICT), which have changed our lives in the past decades. We will relate their lives for just a relatively short period of time. The last 25 years will be enough: from the late eighties, the time Greenaway’s film premiered, until today. Do not think that the experiences of our characters are constrained to a particular country. Their adventures can be placed in the US or Europe, Australia or Canada. Their stories are very similar in any developed country. We will not dare to say identical, but similar, at least broadly¹.

How have the job opportunities of our three characters changed from the end of the eighties? Broadly, job opportunities for the maid and the doctor have increased, while the clerk has been left with fewer opportunities. This means that, in the past 25 years, job distribution has undergone a process of polarisation: both the share of employment in high-skill, high-wage occupations and low-skill, low-wage occupations increased. Job polarisation happened both in the US and Europe (see Figure 1)².

Broadly speaking, the rise in wage inequality from the 1980s was seen only in the US, UK, Canada, and Australia³. There are a range of studies comparing earnings inequalities across OECD and European countries that reveal large country-specific differences in the level of inequality and its increase over time⁴. However, wage structures were roughly stable in Continental European countries, although there is evidence


1. For a broad view, see the academic papers by Acemoglu and Autor (2011), Van Reenen (2011) and Autor (2015).

2. See Autor et al. (2006), Acemoglu and Autor (2011) and Goos, Manning and Salomons (2009, 2010, 2014).

3. See Acemoglu and Autor (2011) and Autor (2015) for the US, Machin and Van Reenen (2007) for the UK, Green and Sand (2015) for Canada and Atkinson and Leigh (2007) for Australia.

4. See, for instance, Burkhauser and Poupore (1997), Gottschalk and Smeeding (2000), Aaberge et al. (2002), Hofer and Weber (2002), Cardoso (2006), Koeniger et al. (2007), Cholezas and Tsakloglou (2007), Checchi and Garcia-Peñalosa (2008), Checchi et al. (2010), Sologon and O’Donoghue (2012), Van Kerm and Pi Alperin (2013) and Bachman et al. (2016).

Figure 1



Source: Acemoglu and Autor (2011). Data on EU employment are from Goos, Manning and Salomons, 2009a. US data are from the May/ORG CPS files for earnings years 1993-2006. The data include all persons ages 16-64, who reported having worked last year, excluding those employed by the military and in agricultural occupations. Occupations are first converted from their respective scheme into 326 occupation groups consistent over the given time period. These occupations are then grouped into three broad categories by wage level."

of widening wage structures starting after the middle of the nineties⁵. Table 1 displays the decile ratios of gross earnings in 11 OECD countries. At the same time, Continental European countries did have a larger increase in unemployment, which may be due to the same underlying forces that have pushed up wage inequality in Britain and America. Differences in institutions, tastes and social norms might explain different cross-country patterns of change, in particular, the differences between Continental European countries and the US and UK⁶. Therefore, in Continental European countries, the maid and the clerk

5. In particular, Dustmann et al. (2009) find that, in Germany, the wages of workers in the lower percentiles of wage distribution fell since the middle of the nineties, while the wages of the upper percentiles increased more than the median. See also Antonczyk et al. (2010).

6. The basic idea is that, in more flexible labour markets – such as the US and the UK – technological shocks result in a price adjustment, and thus wage inequality is observed; while in more rigid labor markets – such as the Continental European countries – technological shocks result in a quantity adjustment, and thus high unemployment is observed.

are not seen to increase their wage difference relative to the doctor – at least not as much as in Anglo-Saxon countries – but maybe at the expense of a higher risk of unemployment.

Table 1

Decile ratios of gross earnings in 11 OECD countries

		1970	1975	1980	1985	1990	1995	2000	2005	2010	2015
Finland	Decile 5/Decile 1			1.49	1.47(d)	1.50	1.41	1.41	1.42	1.45	1.46(a)
	Decile 9/Decile 1			2.47	2.48(d)	2.56	2.34	2.41	2.49	2.52	2.57(a)
France	Decile 5/Decile 1						1.56	1.53	1.48	1.50	1.49(f)
	Decile 9/Decile 1						3.12	3.08	2.97	2.97	2.98(f)
Germany	Decile 5/Decile 1					1.96(g)	1.66	1.71	1.83	1.87	1.87(a)
	Decile 9/Decile 1					3.43(g)	2.90	3.06	3.15	3.33	3.41(a)
Italy	Decile 5/Decile 1				1.56(d)	1.54(l)	1.50	1.44	1.50(j)	1.44	1.50
	Decile 9/Decile 1				2.22(d)	2.54(l)	2.33	2.22	2.40(j)	2.12	2.17
Netherlands	Decile 5/Decile 1							1.59(k)	1.62(l)	1.66	
	Decile 9/Decile 1							2.78(k)	2.91(l)	2.94	
Norway	Decile 5/Decile 1						1.38(m)	1.41	1.46	1.56	1.68
	Decile 9/Decile 1						1.94(m)	2.00	2.13	2.30	2.55
Portugal	Decile 5/Decile 1								1.61	1.44	1.56(a)
	Decile 9/Decile 1								4.31	3.69	3.89(a)
Spain	Decile 5/Decile 1								1.67	1.65	1.64(f)
	Decile 9/Decile 1								3.47	3.30	3.08(f)
Sweden	Decile 5/Decile 1		1.38	1.28	1.28	1.21	1.39	1.39	1.35	1.38	1.36(b)
	Decile 9/Decile 1		2.15	1.96	1.98	1.99	2.21	2.35	2.23	2.23	2.28(b)
United Kingdom	Decile 5/Decile 1	2.08	1.83	1.79	1.83	1.88	1.89	1.84	1.83	1.81	1.78
	Decile 9/Decile 1	3.57	3.02	2.99	3.20	3.43	3.49	3.46	3.64	3.58	3.49
United States	Decile 5/Decile 1		1.92	1.95	2.01	2.03	2.11	2.05	2.10	2.12	2.10
	Decile 9/Decile 1		3.75	3.83	4.13	4.34	4.59	4.49	4.86	5.01	5.04

Source: OECD Database

Much of the increase in inequality in income distribution in many developed countries is explained by the increase in the top 1% ⁷. In English-speaking countries, the income share of the top 1% of income distribution has soared since the late eighties (see Figure 2).

At this stage of the 21st century, in countries like the US or the UK, the top 1% income share is hovering at the figures observed at the beginning of the past century. However, in Continental European countries, the top income shares remained much more stable (see Figure 3 and Figure 4). Figure 5 compares the income share of top 0.1% in the US

7. See Atkinson et al. (2011).

and France, while Figure 6 shows that the top 1% share has increased in many countries, although the rise has been larger in some countries than in others. The largest rise has been reported in the US or Canada, but substantial increases are also seen elsewhere. France exhibits smaller, but still noticeable, increases, from 7.4% to 9%. A substantial part of the rise in US top income inequality represents a rise in labour income inequality, particularly if one includes 'business income' (i.e. profits from sole proprietorships, partnerships and S-corporations) in the labour income category (see Figure 7).

Finally, we can take a wider perspective by focusing on the evolution of the world's income distribution in the past thirty years⁸. Figure 8 shows real income gains obtained at different percentiles of global income distribution between 1998 and 2008. People around the global median (point A) and those who are part of the global top 1% (point C) obtained large real income gains, while real income of people around the 80-85th percentile of the global distribution (point B) did not grow. Who are the people in points A, B, and C? Nine out of 10 people around the world median (i.e. around A) are from Asian countries, mostly from China and India. People at the global top 1% (point C) are mostly workers in the upper halves of the rich countries' income distributions. Finally, people around the 75-90th percentile (point B) are mostly workers in the lower halves of the rich countries' income distributions.

Therefore, from a global perspective, the fate of the maid and the clerk (i.e., low-skill, low-wage workers and middle-skill, middle-wage workers) in rich countries has not been very favourable, while the fate of the doctor (i.e., high-skill, high-wage workers) in rich countries and the fate of workers in poor countries has been promising.

The main reason for this stems from the rise and development of information and communications technologies, which has brought mixed fortunes for them. The doctor has taken advantage of this technological revolution because new technologies complement her work and make her more productive. Moreover, the global markets generated by new technologies have given some exceptional workers the

8. See Milanovic (2016a, 2016b) and Lakner and Milanovic (2015).

Figure 2




Figure 3






Figure 4



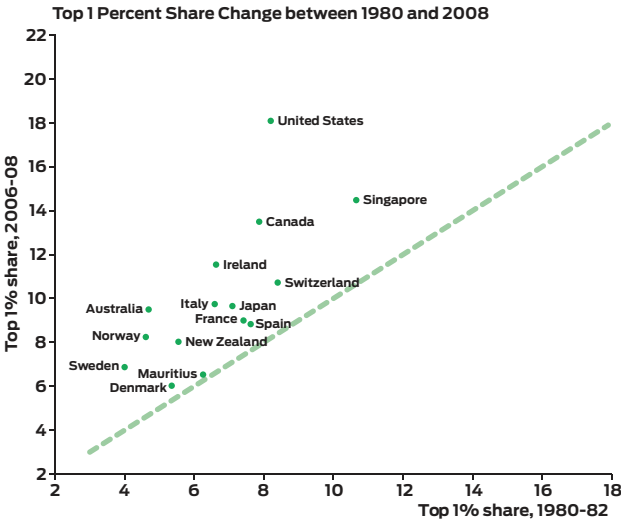
Source, FIGURE 2, 3, 4: Atkinson, Piketty, and Saez (2011).

Figure 5




Source: Source: Jones and Kim (2014).

Figure 6



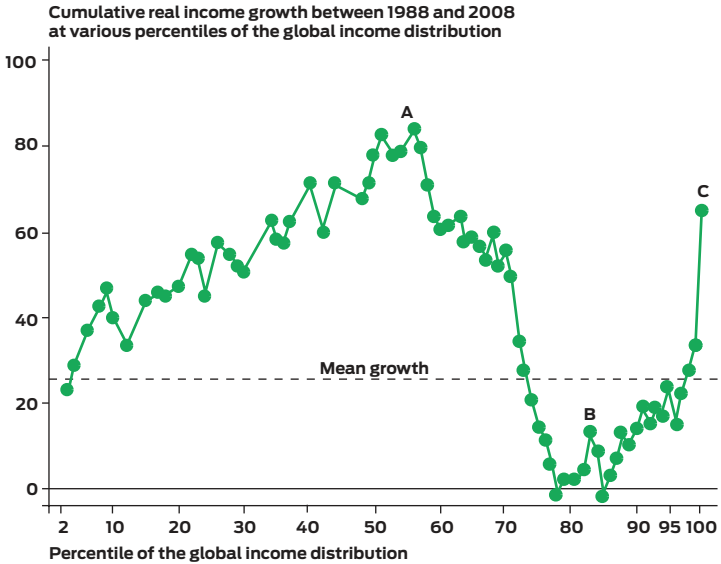
Source: Jones and Kim (2014) Note: Top income inequality has increased since 1980 in most countries for which we have data. The size of the increase varies substantially, however. Data from World Top Incomes Database.

Figure 7



Source: Jones and Kim (2014) Note: The figure shows the composition of the top 0.1 percent income share. These data are taken from the "data-Fig4B" tab of the September 2013 update of the spreadsheet appendix to Piketty and Saez (2003).

Figure 8



Source: Milanovic (2016b)

opportunity of enormous earnings. However, the maid and the clerk, especially the latter, have witnessed how many of their tasks are now undertaken by computers or were relocated to remote places that are no longer so remote thanks to information and communications technologies. Workers in poor countries that have joined the world market have also benefited from globalisation.

Next, we present different and complementary explanations of these facts: automation, offshoring and the existence of superstars.

Automation

The *routinisation hypothesis*, Autor et al. (2003), states that computer and computer-controlled equipment have substituted workers in routine tasks because computers and computer-controlled equipment are highly productive and reliable at performing the tasks that programmers can script, and that are procedural, rule-based activities⁹. This means that the development of information and communications technology has led to a race between machines and workers¹⁰.

Occupations can be classified in four categories according to the kind of performed tasks¹¹:

- i. Non-routine cognitive task-intensive occupations (managerial, professional and technical occupations);
- ii. Routine cognitive task-intensive occupations (sales, clerical and administrative support occupations);
- iii. Routine manual task-intensive occupations (production, craft, repair, and operative occupations); and
- iv. Non-routine manual task-intensive occupations (service occupations, such as food preparation and serving, cleaning and janitorial work, grounds cleaning and maintenance, in-person health assistance by home aides, and numerous jobs in security and protective services).

Non-routine cognitive tasks (i) are highly complementary to information and communications technologies, while non-routine manual tasks (iv), are less complementary to information and communications technologies, since they cannot be replaced by computers and comput-

9. Many examples of routine tasks are given by Levy and Murnane (2004).

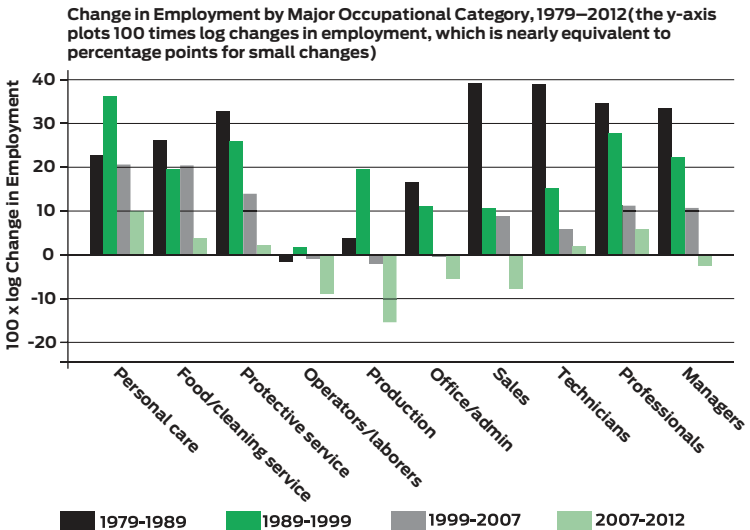
10. The issue is extensively addressed by Brynjolfsson and McAfee in two books (Brynjolfsson and McAfee 2011, 2014). An excellent press article by Autor and Dorn (2013) also addresses this issue.

11. See Acemoglu and Autor (2011) and Autor et al. (2003).

er-controlled equipment because they require situational adaptability, visual and language recognition, and in-person interactions. However, routine tasks – both cognitive (ii) and manual (iii) – can be easily replaced by computers and computer-controlled equipment.

Since the late eighties, the U.S. labour market has undergone: (1) an increase of employment in non-routine cognitive task-intensive occupations (i); (2) an increase of employment in non-routine manual task-intensive service occupations (iv); and (3) a decline of employment in middle-skill, routine task-intensive occupations - (ii) and (iii) - (see Figure 9)¹².

Figure 9




Sources: Autor (2015). Author using data from the 1980, 1990, and 2000 Census IPUMS files, American Community Survey combined file 2006–2008, and American Community Survey 2012. The sample includes the working-age (16–64) civilian noninstitutionalized population. Employment is measured as full-time equivalent workers. **Notes:** Figure 2 plots percentage point changes in employment (more precisely, the figure plots 100 times log changes in employment, which is close to equivalent to percentage points for small changes) by decade for the years 1979–2012 for ten major occupational groups encompassing all of US nonagricultural employment. Agricultural occupations comprise no more than 2.2 percent of employment in this time interval, so this omission has a negligible effect.

12. See Acemoglu and Autor (2011).

The evolution of job opportunities by occupations has been broadly similar in Europe and the US. Figure 10 and Figure 11 show changes in young-male and young-female employment for different occupations in several European countries¹³.

Figure 10

Change in Employment Shares of Young Male Workers (Age<40) by Country 1992-2008



Source: Acemoglu and Autor (2011). European data from Eurostat data 1992-2008. See note to Figure 15. Employment shares are calculated for each of the ten European countries individually, for workers under 40 years of age.


The *routinisation hypothesis* illustrates that technological change does not necessarily imply that everyone wins (i.e., a Pareto-improvement): it generates winners and losers. Workers are more likely to benefit directly from automation if they supply tasks that are complemented by automation, but not if they primarily (or exclusively) supply tasks that

13. The demand for high-skill workers in European countries between 1980 and 2004 increased more in the industries which increased relatively more their use of the information and communications technologies, while the demand for middle-skill workers fell rapidly. However, demand for low-skill workers was broadly unaffected by the investment industries made in information and communications technologies. (see Michaels, Natraj and Van Reenen 2010). These findings give empirical support to the *routinisation hypothesis*.

can be substituted by machines (see Autor 2015). In particular, low-skill, low-wage workers and middle-skill, middle-wage workers employed in routine task-intensive occupations are substituted by machines and their job opportunities decrease, while their wages undergo downward pressure. However, high-skill, high-wage workers employed in non-routine abstract task-intensive occupations benefit from information technology via a virtuous combination of strong complementarities between routine and abstract tasks, elastic demand for services provided by non-routine abstract task-intensive occupations, and inelastic labour supply in these occupations over the short and medium term. These same synergies do not apply to jobs that are intensive in non-routine manual tasks such as janitors, security guards or home health aides, and which are mainly performed by low-skill, low-wage workers. Overall, non-routine manual task-intensive activities are at best weakly complemented by computerisation, do not benefit from elastic final demand, and face elastic labour supply that tempers demand-induced wage increases.

Figure 11

Change in Employment Shares of Young Female Workers (Age<40) by Country 1992-2008



Source: Acemoglu and Autor (2011).

To summarise, the maid, the clerk and the doctor are affected by computers in different ways. The clerk has been replaced in a lot of tasks by computers, which has pressured demand for her services downward and, hence, her wage, while the tasks performed by the doctor have not been replaced by computers – at least, not yet – and they are strongly complementary to the tasks performed by computers. In this context, the doctor has been able to significantly increase her earnings, taking advantage of information and communications technologies. The tasks performed by the maid have not been replaced by computers, but her tasks are only weakly complemented by computerisation.

Estimates of the share of jobs at risk of automation differ widely. Some authors have estimated that between 30% and 60% of the employees in developed countries can be substituted by machines,¹⁴ while others have estimated more modest figures. In particular, Arntz et al. (2016) find the share of jobs at risk of automation is, on average across OECD countries, 9%, even with some heterogeneity across OECD countries (see Figure 12)¹⁵. Differences in the potential for automation, known as automatibility, between educational levels (see Figure 13) and income levels (see Figure 14) are large: workers of low education and low income confront a higher risk of automation¹⁶.

As pointed out by Autor (2015) – who seeks to avert us from falling prey to the ‘Luddite fallacy’ – the negative consequences of automation must be not exaggerated, because automation also complements labour, raises output in ways that lead to a higher demand for labour and interacts with adjustments in labour supply. Moreover, even the modest Arntz et al. (2016)’s figures should be interpreted with cau-


14. See the seminal work by Frey and Osborne (2013) for the US, Pajarinen et al. (2014) and Pajarinen and Rouvinen (2015) for Finland and Norway, Brzeski and Burk (2015) for Germany, and Bowles (2014) for some European countries.

15. As pointed out by Arntz et al. (2016), the approach followed by Frey and Osborne (2013) likely overstates the share of jobs at risk of automation because occupations usually consist of performing a bundle of tasks not all of which may be easily automatable (Autor 2015). Therefore, the potential for automating entire occupations and workplaces may be much lower than suggested by the approach followed by Frey and Osborne (2013). Moreover, the potential for automation must not be confused with actual employment losses. In particular, the substitution may not be reasonable from an economic point of view, as well as impossible for legal reasons or too difficult due to ethical reasons.

16. See Arntz et al. (2016).

tion. Firstly, their approach still reflects technological capabilities, rather than the actual utilisation of such technologies, which might lead to a further overestimation of job automatibility. Secondly, even if new technologies are increasingly adopted in the economy, the effect on employment prospects depends on whether workplaces adjust to a new division of labor, as workers may increasingly perform tasks that are complementary to new technologies. Thirdly, the approach considers only existing jobs, although new technologies are likely to create also new jobs. Moreover, new technologies may also exert positive effects on labour demand if they raise product demand due to improved competitiveness and a positive effect on workers' incomes. Therefore, their findings suggest that fewer workplaces are likely to be 'at risk' than suspected.

Figure 12



Source: Arntz et al. (2016).

Figure 13





Figure 14



Source, Figure 13 and 14: Arntz et al. (2016). Authors' calculation based on the Survey of Adult Skills (PIAAC) (2012)

Offshoring

The development of information and communications technologies has resulted in a drastic reduction in both the cost of processing information and transmitting it between any two points on the globe. Production processes do not necessarily have to be organised locally any more. The possibility of vertical integration has been opened on a planetary scale. This means that the various tasks of a production process can be located in distant parts of the globe, a process known as offshoring.

Offshoring has a significant impact on international trade¹⁷. In particular, trade in intermediate products is gradually becoming increasingly important relative to traditional trade in final products¹⁸; much of the recent surge of international trade comes from middle- and low-income countries; and prices of manufactured imports by developed countries from developing countries are dramatically declining. After the previous stage of globalisation, driven by a sharp reduction in transport costs and political barriers to trade, we are witnessing a second type of globalisation since the late eighties, characterised by what Baldwin (2011) has called connective technologies, which are enabling a drastic reduction in the cost of transmitting information and, thus, facilitating access to knowledge and global connection of people and companies (see Figure 15).

Offshoring induces a direct displacement effect of domestic workers, leading to lower employment and wages, though offshoring activities may also generate a productivity effect similar to technology improvement by lowering a firm's production costs¹⁹. This productivity effect, in turn, will lead to an expansion of output and thus raise employment and wages. The balance between these two forces will determine the direction of the wage and employment effect of offshoring. Therefore, as automation, the main question implied by offshoring is not the end of employment, but that it does not necessarily imply that everybody wins (i.e. a Pareto-improvement).


17. See Feenstra and Hanson (2003).

18. See Hummels et al. (2001).

19. See Grossman and Rossi-Hansberg (2008).

Figure 15

Trade costs (left, 1870 – 2000) & ICT indicators (right, 1975-2011).



Source: Baldwin (2011). Notes: Trade cost measure (left panel) estimated from gravity equations drawing on Jacks, Meissner, and Novy (2011). ICT indicators (right panel). Data for phone and internet users from WDI; data for internet hosts from www.isc.org.

Essentially, any job that does not need to be done in person (i.e., face-to-face) can ultimately be offshored, regardless of whether its primary tasks are abstract, routine, or manual²⁰. With this in mind, it is not surprising that offshoring – made possible thanks to information and communications technologies – is also contributing to altering the relative demand for different kinds of workers and, consequently, affecting their relative wages. Thus, *offshoring* is affecting the three characters of our story differently. The doctor and the maid perform tasks that cannot be offshored, not yet at least, while the tasks of the clerk are increasingly being offshored.

Occupations threatened by computer replacement and those under threat of relocation are not exactly the same. However, in general, it can be asserted that many occupations performed by middle-skill workers

20. See Blinder and Krueger (2013).

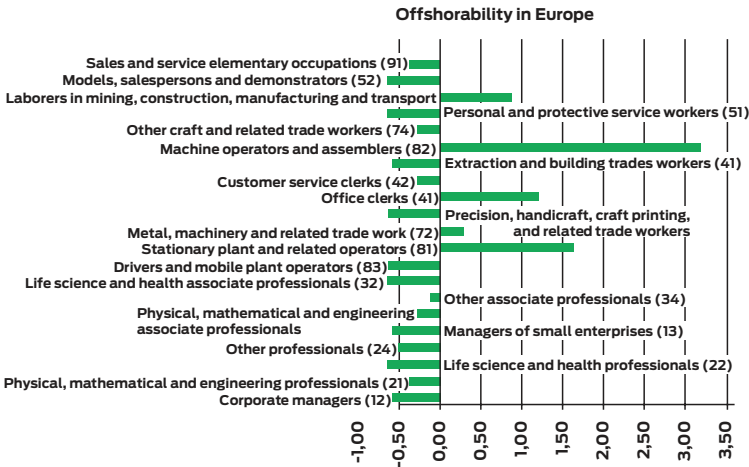
– workers with high-school or some college education – are currently faced by both threats. Acemoglu and Autor (2011) report that offshorability (i) is higher in *clerical and sales occupations*, which imply a high degree of routine tasks; (ii) is also high in *professional, managerial and technical occupations*, which imply a high degree of non-routine tasks, and (iii) is relatively low in *production and operative occupations* and even lower in *service occupations*. According to the estimates provided by Blinder and Krueger (2013)²¹, *farming, fishing and forestry occupations, construction and extraction occupations, installation, maintenance and repair occupations, and transporting and material-moving occupations* have a very low degree of offshorability, but *other production occupations* have a very high degree of offshorability. In particular, Blinder and Krueger estimate that 80% of jobs in *other production occupations* are offshorable, 41% of *office and administrative support occupations*, 17.8% of *sales and related occupations*, 20.5% of *professional and related occupations*, and 16.4% of *management, business, and financial occupations*. However, the percentage of offshorable jobs in *service occupations* is only 0.7%. Blinder and Krueger estimate that roughly 25% of US jobs are offshorable. Figure 16 displays a measure of offshorability of several occupations in 16 European countries elaborated by Goos et al. (2010) (a higher figure means higher offshorability).

Empirically, Autor et al. (2015) report that the US local labour markets, the initial industry composition of which exposes them to rising Chinese import competition, are experiencing significant declines in employment, particularly in manufacturing and among non-college workers. Ebenstein et al. (2014) present evidence that globalisation has put downward pressure on worker wages through the reallocation of workers away from higher-wage manufacturing jobs into other sectors and other occupations. Using a panel of workers, they find that occupation-switching due to trade led to real wage losses of 12 to 17 percentage points. Baumgarten et al. (2010), using a large sample of individual data for Germany, find substantial negative wage effects from offshoring for

21. The measures of offshorability reported by Acemoglu and Autor (2011) are roughly consistent with the measures reported by Blinder and Krueger (2013). However, it must be pointed out that there is a large heterogeneity in the category *production and operative occupations*.

low- and medium-skill workers. However, they also find that the magnitude of these effects strongly depends on the type of tasks workers perform. For instance, they find that, for low-skill workers carrying out tasks with the lowest degree of interactivity, increased offshoring between 1991 and 2006 accounts for a cumulative yearly wage reduction of 1,965 euros, while, for low-skill workers with the highest degree of interactivity, offshoring can only explain a yearly wage reduction of 435 euros. Using Danish data, Hummels et al. (2014) find that offshoring has considerably different wage effects across educational groups, raising skilled labour wages and lowering unskilled labour wages. Moreover, they find that, conditional on skill type, routine tasks suffer wage losses from offshoring and occupations that intensively employ knowledge sets from maths, social science, and languages gain from offshoring shocks, while those that employ knowledge sets from natural sciences and engineering are no more or less insulated from offshoring shocks than the average manufacturing worker.

Figure 16



Note: Rescaled to mean 0 and standard deviation 1, a higher value means more offshorable. Values for ISCO 12 and 13 have been made the same by taking the mean weighted by hours worked.
Source: Goos, M., Manning, A., and A. Salomons. (2010, Table 4, Column 5).

Superstars

So far, our history has been about ordinary workers. There are thousands of maids, clerks and doctors. However, a small number of workers excel. They are *superstars*, like the soccer players Leo Messi or Zlatan Ibrahimovic. In recent decades, there has been an impressive burst of earnings of this type of worker. Recent rises in top earnings have been attributed to the ‘superstar’ phenomenon.

According to the superstar theory, people differ in their talents and the top performer in a field is able to extract payment that is proportionate to the extent of the market served: those unable to afford to see the superstar go to the next-best performer, and so on²². This means that the earnings of the second-best performer depend on the ‘reach’ of the top performer, and so on down the range of talent. That reach has been extended by technology.

Think of the extremely high remuneration for most CEOs of large companies (see Figure 17). Before the eighties, their earnings were certainly enviable, but did not reach, and were not even close to, stratospheric current figures²³. US CEOs are paid significantly more than their European counterparts (see Table 2). Moreover, US executives receive a greater share of their compensation in the form of stock options, restricted shares and performance-based bonuses²⁴. The rise in top-level inequality occurs across a range of occupations; it is not just in finance or among CEOs, for example, but includes doctors and lawyers and star athletes as well²⁵.

To understand the relationship between the earnings of superstars and technological change, consider Luis Suárez (whom we’ll refer to as Suárez I), an F.C. Barcelona footballer in the fifties, and Luis Suárez, the contemporary F.C. Barcelona player (whom we’ll refer to as Suárez

22. The theory of superstars is due to Sherwin Rosen (1981), although he stressed his debt to Alfred Marshall writing many years earlier.

23. See Frydman and Jenter (2010). Gabaix and Landier (2006) give an explanation of the increase in CEO earnings based on the ideas developed below.

24. See Conyon et al. (2011).

25. See Bakija et al. (2012) and Kaplan and Rauh (2010).


II)²⁶. In the fifties and sixties, Suárez I, as does Suárez II today, produced non-rival goods – shots, dribbling, passing, goals – which could be enjoyed by many people at the same time. People paid to see Suárez I performing on the pitch then and now pay to see Suárez II scoring goals. Suárez I, like Suárez II today, earned a lot of money because a lot of people watched and admired his sportsmanship. However, in the fifties, when Suárez I played soccer, technology only allowed 100,000 people to watch Suárez I's games: those were the people who went each Sunday to the Camp Nou Stadium. Nowadays, millions of people around the globe watch Suárez II scoring every week on their televisions, and millions of other people may dream of being Suárez II, playing on their gaming consoles. More people can currently enjoy Suárez II's skills than Suárez I's in the fifties and sixties, thanks to information and communications technologies. Consequently, Suárez II's earnings are much higher than Suárez I ever dreamed.

Moreover, the difference between Suárez II's earnings and most players in the Spanish soccer league is greater than the difference between Suárez I's earnings and most players in the Spanish and Italian leagues in those days. In the fifties, football fans in Vigo could only attend Balaídos Stadium to watch R.C. Celta de Vigo matches, because they could not see to play Suárez I every weekend. They paid for their tickets, as did the supporters of F.C. Barcelona to see games at the Camp Nou Stadium. In the fifties, Barcelona soccer players earned more money than Celta players because more people went to watch them, but not much more. Nowadays, people from Vigo do not need to attend Balaídos Stadium to watch soccer. They can watch F.C. Barcelona matches using pay-per-view digital television services and, consequently, Barcelona soccer players earn much more money than their colleagues at R. C. Celta. This increase in earning difference is a consequence of the 'winner takes all' principle when people produce non-rival goods. Yet, in our example, watching soccer at the stadium is not the

26. The supporters of Deportivo de La Coruña – Suárez I's hometown – were able to enjoy his skills prior to his transfer to F.C. Barcelona. Luis Suárez then won two European Cups with F.C. Internazionale Milano. Luis Suárez is the only Spanish soccer player (actually Galician) to win a Golden Ball. The homonymous Uruguayan player who currently plays in the ranks of F.C. Barcelona was so christened in his honor.

same as watching soccer on TV, watching your team is not the same as watching any other team. There are several reasons why R.C. Celta and F.C Barcelona are not perfect substitutes for each other, meaning the core idea of ‘winner takes all’ does not work perfectly.

Figure 17



Source: Frydman and Jenter (2010). The structure of CEO compensation from 1936 to 2005. The diagram shows the median level and the average composition of CEO pay in the 50 largest firms in 1940, 1960, and 1990 (for a total of 101 firms). Compensation data are hand-collected from proxy statements for all available years from 1936 to 1992; the S&P ExecuComp database is used to extend the data to 2005 (Frydman & Saks 2010). The figure depicts total compensation and the three main components that can be separately tracked over the sample period: salaries and current bonuses, payouts from long-term incentive plans (including the value of restricted stock), and the grant-date values of option grants (calculated using Black-Scholes). The component percentages are calculated by computing the percentages for average CEO pay in each period and then applying them to median CEO pay, as in Murphy (1999). All dollar values are in inflation-adjusted 2000 dollars. Note that the vertical axis is on a log scale.

In short, the development of information and communications technologies is increasing the degree of non-rivalry of some goods produced by some economic agents, which allows them to earn a lot of money if they are the winners of a ‘winner-takes-all’ race. This means that globalisation driven by information and communications technologies has played an important role in boosting demand for the servic-

es of some superstar workers, which has increased their earnings, and hence income inequality²⁷.

Table 2

Summary Statistics the Level and Structure of 2008 CEO Compensation, by Country

Group	Sample Firms	Total Pay		Average Composition of Total Pay			
		Average (000s)	Median (000s)	Base Salary	All Bonuses	Equity Pay	Other Pay
Belgium	28	1,328	884	64%	20%	6%	10%
France	156	1,522	822	60%	21%	15%	4%
Germany	80	2,606	1,739	39%	42%	9%	11%
Ireland	23	2,585	1,375	54%	9%	23%	15%
Italy	46	2,717	2,183	53%	19%	13%	15%
Netherlands	60	1,526	1,166	49%	21%	17%	13%
Sweden	511	,273	1,055	61%	16%	1%	22%
Switzerland	29	3,636	1,336	57%	17%	12%	14%
United Kingdom	419	2,016	1,183	46%	18%	28%	9%
All Europe	892	1,989	1,200	50%	21%	19%	10%
United States	1,426	3,784	2,414	29%	20%	46%	6%

Note: European data from Boardex and US data from ExecuComp exclude firms with less than €100m in 2008 revenues. CEOs in their first year are excluded. Total compensation defined as the sum of salaries, bonuses, benefits, and grant-date values for stock options, restricted stock, and performance shares. US dollardenominated data are converted to Euros using the 2008 year-end exchange rate (€1 = \$1.3919).

Some challenges

In May 2016, Shanghai hosted the CES ASIA trade fair for the consumer technology industry. A company named Bubble Lab presented a sophisticated robotic arm capable of preparing tea, coffee, and cocktails. After service, it even cleans the table and leaves it spotless. Mr. Shen Li, representing the company, said, “I do not know if it’s true that [it] is going to end human work, but the fact is that machines can perform ever-more complicated tasks”. So far, computers and computer-controlled equipment have replaced human labour in a wide variety of tasks. Yet, as the

27. Haskel et al (2012) suggest that globalization may have raised the returns to superstars via the Rosen (1981)’s mechanism.

above anecdote illustrates, technological developments can widely expand the set of tasks that can be performed by machines. Many already conjecture that all tasks not requiring creativity will be undertaken in the near future by machines.

Such a scenario means we should think seriously about the challenges faced by society. Seven key issues should be the focus of our attention: income redistribution; education and job training; taxes and regulation; the welfare state; convergence and development policies, the trade-off between equity and efficiency; and, political organisation.

Income redistribution

A question arises: in the event that computer equipment replaces most tasks now performed by human labour, what are those who do not have the required skills, namely creativity, waiting for? The US film *Elysium*, released in 2013, describes a dystopian future in which machines perform most tasks. People who have been replaced by machines barely survive on a degraded planet Earth, while a small elite inhabits an artificial satellite orbiting our planet. The film is just science fiction. However, you can imagine a society in which a high-skill minority performs tasks highly complementary to technological capital and around which most of income and wealth is concentrated, while most people survive by producing low-priced goods in a highly automated production process²⁸.

We do claim we will get to this situation. Nonetheless, in a future in which the automation of both production and increasingly important non-rival goods give rise to a polarised society, we must consider new and more suitable ways for income redistribution. Proposals such as a basic income²⁹, the distribution of ownership of some production assets³⁰, or a negative income tax³¹ are being discussed, especially the latter. The big challenge comes from combining new mechanisms of re-

28. See Cowen (2013).

29. See Van Parijs (1995).

30. See Paine (1797) and Roemer (1994).

31. See Friedman (1968).

distribution with the necessary incentives for prosperity. However, the advantages of these redistributive proposals should not be ignored, because they would enable direct state intervention to be reduced in the provision of goods and services, such as education, health, social assistance and social insurance.

However, polarisation might be transitory because, while some of the tasks in many current middle-skill jobs are susceptible to automation, many middle-skill jobs will continue to demand a mixture of routine and non-routine tasks from across the skill spectrum. Thus, a significant stratum of middle-skill jobs combining specific vocational skills with foundational middle-skills levels of literacy, numeracy, adaptability, problem solving and common sense, will endure in the coming decades³².

Education and job training

The education and job training system must prepare for the kinds of workers that will thrive in these middle-skill jobs of the future. Therefore, the education system must produce skills that complement rather than are substituted by, technological change.

An education system that fosters creativity and provides the suitable knowledge for the new technological reality is urgently needed in order to successfully meet new technological challenges. We agree that this is nothing but a vague generalisation. However, many authors from different fields have recognised that educational systems suffer from sclerosis. A well-known education advisor, has criticised educational systems, and asserted that (nowadays) schools kill creativity³³. Goldin and Katz (2008) blame the US educational system for not having sufficiently adapted to the demands of the new reality and they argue it is responsible, at least partially, for the increase in wage inequality in the US. Brynjolfsson and McAfee (2014) recommend redirecting the educational system from its focus on reading and mathematics, typical of the industrial era, towards a broader set of intellectual and personal

32. See Autor (2015).

33. See Robinson (2007, 2015).

skills. Even UNESCO sponsored a World Education Forum “Rethinking Education” in 2015. Thus, this institution is aware that something is going wrong in educational systems worldwide.

We do not want to be so presumptuous as to know what changes the education system exactly needs. Nonetheless, our conviction is that the inability of the education system to adapt to these new times has underpinned state interventionism. Overcoming atrophy requires that state interventionism in education decreases. Thereby, new educational alternatives will arise from competition and the subsequent creative destruction. Introducing school vouchers would be helpful³⁴. Vouchers already exist in Sweden and Denmark, but not in most European countries. However, a school voucher system is not enough. The state must give up its tight control of the education system. Educational innovation will not be possible if it is subject to restrictions imposed by a straitjacket.

In a changing world, in which many workers are at risk of being displaced from their jobs by machines, it is important to be very aware of the need to recycle these workers. Therefore, the reskilling of workers and training support for the unemployed are key issues. Once again, the best way to achieve objectives in this area is to abandon direct state intervention and to allow markets to work. To this end, severance pay should be substituted by periodic contributions to a worker’s capitalisation funds (the so-called Austrian fund), which could be used for reskilling, and unemployed workers should be provided with training vouchers to finance their preferred training courses, given by the provider that they consider most suitable.

Taxes and regulation

Tax systems should also be adapted to the new technological reality. Spain is a good example, with labour incomes excessively burdened by taxes. If machines and overseas labour are replacing domestic workers in many tasks, labour income can no longer bear the tax burden. In particular, welfare state funding should no longer mainly fall on labour in-

34. See Friedman (1955).

come, as currently happens in many countries. Contributions to social security and other charges should be cut, and more flexible labour relationships should be allowed.

Market regulation could hinder innovation. Europe's regulation of GMOs, or the difficulties faced by the so-called sharing economy to entering hyper-regulated markets, are some examples of how regulation can hinder innovation. Particular consideration should be given to financial regulation. The regulation of financial markets should also be reduced. Otherwise, innovation may suffer due to difficulties in finding funding. Thus, eliminating some regulatory barriers and rethinking regulation seems necessary to face the new technological reality.

Moreover, polarisation in the labour market reflects large sections of the middle class may currently be adversely affected by new technologies, at least in the short term. Everyone is aware of the importance of middle-class preferences in determining public policy in a democratic society. The reaction to technological change may lead, therefore, to successful demands for setting higher barriers to trade or to technological adoption, as well as to pressure to implement redistributive policies in favour of these sections of the middle-class (more public employment, for instance). One cannot help thinking that some recent political outcomes in Europe have followed this logic.

Welfare state

Reducing the tax burden on labour without a dramatic increase in other taxes might be only possible if the efficiency of the welfare state is improved. This would likely require new organisational models in public services. Greater individual freedom to choose and increased competition would help to promote efficiency. A higher weighting of private sector involvement in providing health services and health insurance and higher private participation in employment insurance and active employment policies are some possible directions³⁵.

Funding a pay-as-you-go 'Bismarck-style' pension system, which is common in continental Europe, is very difficult in an extremely polar-

35. The Dutch health system, one of the best in Europe, could be a reference model.

ised labour market. Moreover, it would give rise to great inequalities in pensions, because in a pay-as-you-go pension system the pension received is proportional to the contribution. In a polarised society, a pay-as-you-go ‘Beveridge-style’ pension system, which is common in Anglo-Saxon countries, financed in the most part by taxes, may be a more suitable model because it would mean a lower proportion of GDP is devoted to the public pension system and because it would give rise to a lower inequality in pensions³⁶. In another direction, building on the ‘Austrian fund’, would give rise to a more-capitalised pension system, accompanied by a higher level of private participation in its management.

Moreover, substituting machines for workers and the growing importance of new forms of human capital might affect the birth rate, which would have a significant long-run impact on pay-as-you-go pension systems³⁷. Additionally, it is also foreseeable that technical advances will prolong the lives of people, which will also negatively affect financial sustainability of pay-as-you-go pension systems. Moving towards a Beveridge pension system and higher capitalisation of the pension system will help to guarantee its solvency.

Convergence and development policies

Information and communications technologies are expanding and facilitating access to information (and, thus, to knowledge and technology) to everyone anywhere. This, together with reduced transport costs and the elimination of political obstacles to mobility of goods and production, is facilitating interactions at a global scale and technological adoption by developing countries. Friedman (2006) asserts that, after the end of the eighties, we are in a new stage of globalisation (which he calls globalisation 3.0).

36. Usually Anglo-Saxon countries devote a lower fraction of their GDP to funding the pension system than Continental European countries. The main reason is that the spirit of a pension system à la Beveridge is to provide all individuals a basic pension.

37. The long-run return rate of a pay-as-you-go pension system equals the population growth rate plus the productivity growth rate.

The best development policy in a global world in which information freely flows is to spread prosperity-promoting institutions. If technological progress is equalising opportunities for countries (flattening the world, as Friedman says), then one should not be surprised to observe in the near future the proliferation of spectacular economic miracles and a rapid change in the geo-economic map. However, the spread of prosperity-promoting institutions around the world cannot stop. Institutions securing property rights and ensuring a free and open society are indispensable for innovation and accumulation. If so, access to information facilitates technological adoption and, consequently, the rapid convergence of the laggards to the most advanced will be seen. However, the flip side must be borne in mind. A country rests on its laurels may experience a rapid relative decline.

The trade-off between equity and efficiency

Inequality has increased less in continental Europe than in the US, which may have positive aspects. However, it can also be a sign of a trade-off between efficiency and equity. Perhaps the egalitarian policies of European countries are preventing them from taking advantage of all the benefits that technological progress and globalisation are putting at their fingertips. In this sense, economic liberalisation can lead to greater inequality, but also to greater innovation and growth.

On one hand, innovation, to the extent that it means the efforts of fast-growing entrepreneurs to improving their products or increasing the productivity of their effort, can increase inequality at the top end of the scale. Yet innovation enhances creative destruction, which can decrease this top-end inequality³⁸. Therefore, the desirable institutional context for innovation is one in which regulation and other public policies do not prevent creative destruction and, consequently, low-income new entrants substituting high-income incumbents.

38. See Jones and Kim (2014).

Political organisation

Finally, the main problem that any organisation must solve is how to transmit all relevant information to all concerned agents. Information and communications technologies facilitate the processing and transmission of information. Therefore, in the near future, the reorganisation of the representative organisation of societies seems to be inevitable under the new information and communications technologies scenario. These technologies will change the relationship between the citizens and the State, with more active participation in social issues of the former and an improvement in the working of the latter. Some attempts, such as e-government, have been headed in this direction. Yet, this process is not without conflicts. More informed and educated citizens' needs might collide with existing corrupt (or even oppressive) regimes, or even extractive political structures within democracies. So, social movements challenging the political and social establishment are likely.

To summarise

We have followed our three characters along the past three decades. During this period, information and communication technologies have increasingly facilitated information processing and transmission, which gives rise to automation and offshoring of many routine tasks. High-skill workers performing non-routine tasks (i.e., the doctor of our story) have taken advantage of their complementarity with the new technologies. They have increased their job opportunities and their earnings. However, low-skill, low-wage workers (i.e., the maid) and middle-skill, middle-wage workers (i.e., the clerk) in rich countries performing routine tasks have been substituted by machines, while a lot of these tasks have been relocated to poorer countries. These trends have deteriorated their job opportunities and caused the obsolescence of their knowledge, particularly true for the clerk. Moreover, globalisation, which was made possible by new information and communication technology, has given some workers the opportunity to earn enormous benefits, resulting in a large increase in the income share of the top 1%.

These processes, originated by recent technological change, are leading to job polarisation and income inequality in rich countries. Thus, globalisation has benefited the standard of living of most of the population in the World, but has had a less-positive impact on – and even hampered – that of lower-skill, lower-wage population in developed countries.

The process also opens new challenges to democratic societies, with wide segments of the population facing the risk of impoverishment, that could call for populist political solutions. Income redistribution, improving educational systems and job training, reformulating the welfare state, tax systems and regulation are some proposals we have suggested to mitigate these challenges and to improve the wellbeing and the economic possibilities of all citizens in modern democratic societies.

References

- Aaberge, R., A. Björklund, M. Jäntti, M. Palme, P. J. Pedersen, N. Smith, and T. Wennemo.** 2002. 'Income Inequality and Income Mobility in the Scandinavian Countries Compared to the United States'. *Review of Income and Wealth*, 48(4), 443–69.
- Acemoglu, D., & Autor, D. H.** 2011. 'Skills, tasks and technologies: Implications for employment and earnings'. *Handbook of Labor Economics*, 4, 1043–1171.
- Antonczyk, D., DeLeire, T., & Fitzenberger, B.** 2010. 'Polarization and rising wage inequality: comparing the US and Germany'. ZEW-Centre for European Economic Research Discussion Paper, (10-015).
- Arntz, M., T. Gregory and U. Zierahn** (2016). 'The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis'. *OECD Social, Employment and Migration Working Papers* 189.
- Atkinson, A. B., and A. Leigh,** 2007. 'The distribution of top incomes in Australia'. *Economic Record*, 83(262), 247–261.
- Atkinson, A. B., Piketty, T., & Saez, E.** 2011. 'Top incomes in the long run of history'. *Journal of Economic Literature*, 49(1), 3–71.
- Autor, D. H.** (2015). 'Why Are There Still So Many Jobs? The History and Future of Workplace Automation'. *Journal of Economic Perspectives*, 29(3), 3–30.
- Autor, D. H., & Dorn, D.** 2013. 'How technology wrecks the middle class'. *The New York Times*, 24.
- Autor, D. H., Dorn, D., & Hanson, G. H.** 2015. 'Untangling trade and technology: Evidence from local labour markets'. *The Economic Journal*, 125(584), 621–646.
- Autor, D., L. Katz, and M. Kearney** 2006. 'The polarization of the US labor market'. *American Economic Review*, 96, 189–194.
- Autor, D., Levy, F., & Murnane, R. J.** 2003. 'The skill content of recent technological change: An empirical exploration'. *The Quarterly Journal of Economics*, 118(4), 1279–1333.

Bachmann, R., P. Bechara and S. Schaffner 2016. 'Wage inequality and wage mobility in Europe'. *Review of Income and Wealth*, 62(1), 181-197.

Bakija, J., A. Cole and B. Heim 2012, 'Jobs and Income Growth of Top Earners and the Causes of Changing Income Inequality: Evidence from U.S. Tax Return Data'. Working Paper Williams College 2010-22

Baldwin, R. 2011. 'Trade and industrialisation after globalisation's 2nd unbundling: How building and joining a supply chain are different and why it matters'. NBER Working Paper 17716.

Baumgarten, D., Geishecker, I. and G. Holger 2010. 'Offshoring, Tasks, and the Skill-Wage Pattern'. IZA Discussion Paper 4828.

Blinder, A., & Krueger, A. B. 2013. 'Measuring offshorability: A survey approach'. *Journal of Labor Economics*, 31(S1), S97-S128.

Bowles, J. 2014, *The Computerization of European Jobs*, Bruegel, Brussels.

Brynjolfsson, E., & McAfee, A. 2011. 'Race against the machine: How the digital revolution is accelerating innovation, driving productivity, and irreversibly transforming employment and the economy'. Digital Frontier Press, Lexington, MA.

Brynjolfsson, E., & McAfee, A. 2014. 'The second machine age: work, progress, and prosperity in a time of brilliant technologies'. WW Norton & Company.

Brzeski, C. and I. Burk 2015. *Die Roboter kommen. Folgen der Automatisierung für den deutschen Arbeitsmarkt [The Robots Come. Consequences of Automation for the German Labour Market]*. ING DiBa Economic Research.

Burkhauser, R. V. and J. G. Poupore 1997. 'A Cross-National Comparison of Permanent Inequality in the United States and Germany', *Review of Economics and Statistics*, 79, 10-17.

Cardoso, A. R. 2006. 'Wage Mobility: Do Institutions Make a Difference?'. *Labour Economics*, 13(3), 387-404.

Checchi, D. and C. Garcia-Peñalosa 2008. 'Labour Market Institutions and Income Inequality'. *Economic Policy*, 23(56), 601–49.

Checchi, D., V. Peragine, and L. Serlenga 2010. 'Fair and Unfair Income Inequalities in Europe'. IZA Discussion Papers 5025, Institute for the Study of Labor (IZA).

Cholezas, I. and P. Tsakoglou 2007. 'Earnings Inequality in Europe: Structure and Patterns of Inter-Temporal Changes'. IZA Discussion Papers 2636, Institute for the Study of Labor (IZA).

Conyon, M., Fernandes, N., Ferreira, M., Matos, P., and K. Murphy 2011. 'The executive compensation controversy: a transatlantic analysis'. ICS 2011-002. Retrieved from Cornell University, ILR

Cowen, T. 2013. 'Average is over: powering America beyond the age of the great stagnation'. Penguin Publishing Group.

Dustmann, C., Ludsteck, J., and Schonberg, U. 2009. 'Revisiting the German Wage Structure'. *The Quarterly Journal of Economics*, 124(2), 843–881.

Ebenstein, A., Harrison, A., McMillan, M., and Phillips, S. 2014. 'Estimating the Impact of Trade and Off-shoring on American Workers Using the Current Population Surveys'. *The Review of Economics and Statistics*, 96(4), 581–595.

Feenstra, R. and G. Hanson 2003. 'Global Production Sharing and Rising Inequality: A Survey of Trade and Wages'. In *Handbook of International Trade*, edited by E. Kwan Choi and James Harrigan, 146–85. Malden, MA: Blackwell Publishing.

Frey, C.B. and M.A. Osborne 2013. *The Future of Employment: How Susceptible are Jobs to Computerization?*. University of Oxford.

Friedman, M. 1955. 'The role of government in education'. *Economics and the Public Interest*.

Friedman, M. (1968). 'The case for the negative income tax: a view from the right'. *Issues in American Public Policy*. Ed. JH Bunzel. Englewood: New Jersey, 111–120.

Friedman, T. L. 2006. 'The world is flat: The globalized world in the twenty-first century' (pp. 3-543). London: Penguin.

Frydman, C., & Jenter, D. 2010. 'CEO Compensation'. *Annual Review of Financial Economics*, 2(1), 75-102.

Gabaix, X. and Landier, A. 2006. 'Why has CEO pay increased so much?' NBER Working Paper 12365.

Goldin, C., & Katz, L. 2008. *The race between technology and education*. Cambridge, MA: Harvard.

Goos, M., Manning, A., and A. Salomons 2009. 'Job Polarization in Europe.' *The American Economic Review, Papers and Proceedings*, 99(2), 58-63

Goos, M., Manning, A., and A. Salomons 2010. 'Explaining job polarization in Europe: the roles of technology, globalization and institutions'. CEP Discussion Paper N° 1026, LSE.

Goos, M., A. Manning, and A. Salomons (2014). 'Explaining Job Polarization: Routine-Biased Technological Change and Offshoring'. *American Economic Review* 104(8): 2509-26.

Gottschalk, P. and T. M. Smeeding 2000. 'Empirical Evidence on Income Inequality in Industrialized Countries'. in A. B. Atkinson and F. Bourguignon (eds), *Handbook of Income Distribution*, North Holland, Amsterdam, 261-307.

Green D. A. and B. M. Sand 2015. 'Has the Canadian Labor Market Polarized?'. *Canadian Journal of Economics*, 48(2), 612-646.

Grossman, G. M. and Rossi-Hansberg, E. 2008. 'Trading Task: A Simple Theory of Offshoring'. *American Economic Review*, 98:1978-1997.

Haskel, J., R. Z. Lawrence, E. E. Leamer and M. J. Slaughter 2012. 'Globalization and U.S. Wages: Modifying Classic Theory to Explain Recent Facts'. *Journal of Economic Perspectives*, 26(2), 119-40.

Hofer, H. and A. Weber 2002. 'Wage Mobility in Austria 1986-1999'. *Labour Economics*, 9(1), 563-77.

Hummels, David, Jun Ishii, and Kei-Mu Yi, 2001. 'The Nature and Growth of Vertical Specialization in World Trade'. *Journal of International Economics*, 54(1), 75-96.

Hummels, D., Jørgensen, R., Munch, J., and C. Xiang, 2014. 'The wage effects of offshoring: Evidence from Danish matched worker-firm data'. *The American Economic Review*, 104(6), 1597-1629.

Jones, C. I. and J. Kim, 2014. 'A Schumpeterian model of top income inequality'. NBER Working Paper 20637.

Kaplan, S. N. and J. Rauh, 2010. 'Wall Street and Main Street: What Contributes to the Rise in the Highest Incomes?' *Review of Financial Studies*, 23 (3), 1004-1050.

Koeniger, W., M. Leonardi, and L. Nunziata, 2007. 'Labour Market Institutions and Wage Inequality'. *Industrial and Labour Relations Review*, 60(3), 340-56.

Lakner, C. and B. Milanovic, 2015. 'Global income distribution: from the fall of the Berlin Wall to the Great Recession'. *World Bank Economic Review* 30(2), 203-232.

Levy, F. and R. J. Murnane, 2004. *The New Division of Labor: How Computers Are Creating the Next Job Market*. Princeton University Press.

Machin, S. and Van Reenen, J. 2007. 'Changes in Wage Inequality. Centre for Economic Performance', Special Paper N^o18.

Michaels, G., Natraj, A., & Van Reenen, J. 2010. 'Has ICT Polarized Skill Demand? Evidence from eleven Countries over 25 Years'. NBER working paper 16138.

Milanovic, B. 2016a. *Global Inequality: A New Approach for the Age of Globalization*. Harvard University Press.

Milanovic, B. 2016b. 'The greatest reshuffle of individual incomes since the Industrial Revolution'. CEPR's Political Portal. <http://voxeu.org/article/greatest-reshuffle-individual-incomes-industrial-revolution>.

Paine, T. 1797. 'Agrarian Justice'. In Cunliffe, J. & Erreygers, G. (eds.) (2004) *The Origins of Universal Grants* (pp. 3-16). Palgrave Macmillan UK. [unbold?]

Pajarinen, M., P. Rouvinen and A. Ekeland, 2014. 'Computerization Threatens One Third of Finnish Employment'. *ETLA Brief*, 22, 1-6.

Pajarinen, M. and P. Rouvinen, 2015. 'Computerization Threatens One Third of Finnish and Norwegian Employment'. *ETLA Brief*, 34, 1-8.

Robinson, K. 2007. Sir Ken Robinson: 'Do schools kill creativity?', TED, released 6 January 2007. <https://www.youtube.com/watch?v=iG9CE55wbtY>

Robinson, K. & Aronica, L. 2015 *Creative Schools: The Grassroots Revolution That's Transforming Education*. Viking.

Roemer, J. E. 1994. *A future for socialism*. Harvard University Press.

Rosen, S. 1981. 'The economics of superstars'. *The American Economic Review*, 71(5), 845-858.

Sologon, D. M., and C. O'Donoghue 2012. 'Equalizing or Disequalizing Lifetime Earnings Differentials? Earnings Mobility in EU: 1994-2001', in J. Bishop (ed), *Inequality, Mobility, and Segregation: Essays in Honor of Jacques Silber*, *Research in Economic Inequality*, Vol. 20, Emerald Group Publishing, 2012.

Van Kerm, P. and M. N. Pi Alperin 2013. 'Inequality, Growth and Mobility: The Intertemporal Distribution of Income in European Countries 2003-2007'. *Economic Modelling*, 35, 931-9.

Van Parijs, P. 1995. 'Real Freedom for All: What (If Anything) Can Justify Capitalism?'. Clarendon Press.

Van Reenen, J. 2011v. 'Wage inequality, technology and trade: 21st century evidence'. *Labour Economics*, 18(6), 730-741.

Georg Graetz is Assistant Professor professor in the Department of Economics at Uppsala University, research affiliate at the Institute for the Study of Labor (IZA), and associate at the Centre for Economic Performance in London. His research focuses on technological change, wage inequality, and the economics of education.

The impact of technological change on the labour market¹

Introduction

How does technological change affect the labour market? Do new machines, devices, and computer software lead to widespread job destruction? How do they change the nature of jobs still done by human workers? What happens to workers' wages, and to overall wealth and living standards? How is the distribution of income affected?

Questions like these have long been studied by economists, but over the past ten years or so there has also been a surge in the public's interest in these issues². An explanation for this surging interest may lie in the coincidence of some surprising technological breakthroughs – such as the driverless car or the IBM Watson supercomputer – with weak labour market performance in most developed countries following the

1. I thank the editors and two anonymous referees for excellent suggestions.

2. See the numerous articles and special reports on automation in major news outlets such as *The New York Times* or *The Economist*, and the success of books such as Erik Brynjolfsson and Andrew McAfee's *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies* (2014) and similar works.

Great Recession³. Whatever the reason, economists can draw on a large body of research to inform the public debate.

This chapter summarises findings from a subset of this research, concerning two particular technologies and their impact on the labour market: namely, information and communication technology (ICT) and industrial robots. While neither of these technologies has led to major job destruction so far, they have changed the kinds of tasks people do, and have had an unequal impact across skill groups. Overall, they have substantially contributed to productivity growth, thus raising overall wealth and living standards.

While the findings described are selective (partly drawing on my own work), they are broadly in line with research on technology and work in general⁴. They may thus give rise to cautious optimism regarding the impact of future innovations. Of course, it may be that “this time is different”, meaning that the past offers little guidance for the future. Indeed, several authors argue that recent innovations have dramatically increased the scope of what machines can do; and have also lowered the transaction costs of conducting business, especially services (Brynjolfsson & McAfee 2014, Ford 2016). The implications could be that a much larger fraction of workers – including highly skilled ones – may be adversely affected by technological change in the future; and that alternative work arrangements, such as independent contracting, will increasingly displace traditional employment relationships, since lower transaction costs allow even one-person ‘firms’ to operate efficiently (Coase 1937). While economists are only beginning to investigate these recent developments, I argue that past patterns of technological change remain an informative guide to policy.

Another threat that some writers point to is the slowdown in productivity growth in recent decades (Gordon 2016; Breman, in this volume, offers an excellent summary of the debate). Clearly, if technological change does not deliver improvements in overall efficiency, then it

3. To be sure, neither the driverless cars nor Watson have been adopted commercially yet, so they cannot be blamed for weak labour markets.

4. For a comprehensive treatment of long-run technological change and its effects on labour markets, see Goldin & Katz (2008). On developments during the past four decades, see Acemoglu (2002) and Acemoglu & Autor (2011).

cannot be true that the ‘winners’ are able to more than compensate the ‘losers’. I believe the slowdown is temporary, as firms need time to figure out how to best apply the latest innovations, a process that can be accelerated by appropriate policies (Breman 2016).

Computers: Polanyi’s paradox and the importance of task complementarity

Over the past two centuries, and especially since the end of World War II, computing performance has increased by a factor of about two trillion, relative to manual calculation, and the price of computing power, relative to the cost of human labour, has decreased by a factor of about 70 trillion (Nordhaus 2007). It is not surprising, then, that technologies that harness this vastly improved computing power, have by now become ubiquitous in most workplaces. But, despite fears of widespread job losses, the number of people employed has steadily increased during the 20th century in most rich countries, in absolute terms and relative to population size⁵.

As overall employment growth remained strong during a phase of fast technological change, the labour market transformed in important ways. Regarding the composition of the labour force, both the proportion of females and college-educated workers increased substantially. At the same time, the distribution of labour income became more unequal. For instance, the wages earned by college-educated workers relative to those with a high-school degree steadily increased. Indeed, the rising return to skill, despite an increase in the supply of skill, appeared to many economists to be one of the most important consequences of the ICT revolution. The dominant explanation was that ICT raised the productivity of high-skill workers disproportionately, meaning they earned relatively higher wages.⁶

Why, and how exactly did computers raise the productivity of skilled workers more than that of the less skilled? In a seminal article, David

5. For a discussion of automation fears and why they have not yet materialised, see Autor (2015).

6. See Acemoglu & Autor (2011) for a review of the large literature on ICT and inequality.

Autor, Frank Levy, and Richard Murnane⁷ pointed out that computers are capable of performing tasks that are well defined in scope, as well as repetitive and hence predictable. Such tasks can easily be codified, that is, instructions for completing the task can be expressed in computer code. Autor and his co-authors called such tasks ‘routine’⁸. They then developed a method of quantifying the routine-ness of several hundred occupations, and documented, first, that industries that initially made greater use of routine tasks adopted computers at a higher rate, and second, that routine occupations, such as accountants, clerks, and some production workers, experienced relative (in some cases absolute) declines in employment.

To understand why computers are limited in the range of tasks that they can perform, Autor and his co-authors appeal to the importance of tacit knowledge as expressed in Michael Polanyi’s famous line “we can know more than we can tell” (Polanyi 1966): while most humans are able to ride a bicycle or can distinguish a cat from a dog based on their appearances, no-one has yet figured out what precisely needs to happen in our brains to accomplish these tasks, and therefore we lack knowledge of the explicit instructions needed for machines to perform them. This insight is known as Polanyi’s paradox. It explains why tasks that require analytical, interactive, and creative skills are not (yet) subject to replacement by computers. But workers performing these ‘non-routine’ tasks may still be affected by ICT, but in a positive way: if routine tasks are critical, or ‘complementary’ inputs for non-routine work, then the rise of ICT makes non-routine workers more productive. For instance, a business consultant who relies on numerical calculations to help with strategic decisions, benefits greatly if these calculations are available at a lower cost. Task complementarity thus helps explain why ICT has made skilled workers disproportionately more productive.

The task framework developed by Autor et al. contained an important implication, as was first pointed out by Maarten Goos and Alan

7. Autor et al. (2003).

8. They use the word differently from its everyday usage. For instance, they consider driving a highly ‘non-routine’ task, given the difficulties in writing computer code to deal with for instance oncoming traffic when taking a left-turn on a busy road. But surveyed truck drivers, applying everyday usage of the term, did describe their work as routine (Autor 2013).

Manning: since many non-routine (that is, non-codifiable) tasks can be performed by untrained workers (think of cleaning, waiting tables, driving), and because many routine occupations were traditionally performed by semi-skilled workers, the ICT revolution should have led to a shift of employment out of middle-wage occupations and towards not only high-wage occupations, but also low-wage ones. Indeed, Goos and Manning did present evidence that such ‘job polarisation’ has occurred in the UK, and later work confirmed the existence of job polarisation in most other European countries and the US⁹.

To sum up, the ICT revolution did not lead to widespread job losses, but it has caused a large decline in the fraction of workers performing tasks that could easily be codified and were thus prone to be automated. Workers are now more likely to perform tasks demanding analytical, interactive, and creative abilities, and their productivity in performing these tasks has received a large boost from improvements in computing power. While the demand for many low-skill jobs has remained steady, workers performing these jobs did not see their wages rise as spectacularly as did high-skill workers. ICT has also contributed substantially to overall productivity growth and thus led to higher living standards, although these gains have not been evenly distributed.

Industrial robots: the workerless factory¹⁰

Since the beginning of the industrial revolution, production has steadily become more automated. The latest and most advanced stage of this process is marked by modern industrial robots. These are autonomous, flexible, and versatile machines, able to perform a wide range of tasks including welding, painting, packaging and others, with very little or no human intervention. As industrial robots are now flexible enough to also connect the various steps of the production chain, some factories can make do without human workers, except for setup

9. See Goos & Manning (2007), Goos et al. (2009), and Autor et al. (2006).

10. This section draws on an article Guy Michaels and I wrote for IZA Newsroom titled ‘Robots at work: Boosting productivity without killing jobs?’, published online on 31 March 2015 and available at <http://newsroom.iza.org/en/2015/03/31/robots-at-work-boosting-productivity-without-killing-jobs/>.

and maintenance tasks. Studying industrial robots potentially offers interesting lessons, because these machines are narrower in scope than ICT, and are by their nature an automating technology, whereas ICT often also involves novel production processes, as well as novel products and services.

In a joint and on-going research project with Guy Michaels¹¹, we compiled a new dataset spanning 14 industries (mainly manufacturing industries, but also agriculture and utilities) in 17 developed countries (encompassing European countries, Australia, South Korea, and the US), which includes a measure of the use of industrial robots employed in each industry, in each of these countries, and how it has changed from 1993-2007. Our data on these robots come from the International Federation of Robotics (IFR). We obtained information on other economic performance indicators from the EUKLEMS database.

The data show a striking increase in the adoption of robots between the early 1990s and the mid-2000s. Robot density – the number of robots per million hours worked – increased by about 150% over this period. This was most likely driven by an equally dramatic fall in prices: quality-adjusted prices of industrial robots fell by about 80%. The rise in robot use was particularly pronounced in Germany, Denmark, and Italy. The industries that increased robot use most rapidly were the producers of transportation equipment, the chemical and metals industries.

When analysing the effects of robot use on employment, productivity, and other outcomes, we face the challenge that other factors may be driving both robot adoption and the outcomes of interest, so that associations between increased robot use and the outcomes may not necessarily be due to a causal effect of robots. To address this and related concerns, we isolate variation in robot adoption that is caused by differences in the replaceability of labour across industries. Such differences are simply due to the fact that the share of labour tasks that can be done by robots (including for instance cutting, welding and painting) is higher in some industries than in others. Our measure of replaceability is constructed based on data from the period prior to widespread robot

11. See Graetz & Michaels (2015) and updates on my website <https://sites.google.com/site/georg-graetz/>.

adoption. The ‘replaceability index’ strongly predicts increased robot use. As an important check on the validity of this exercise, we find no significant relationship between replaceability and the outcomes before the adoption of robots.

How did the increased use of robots affect human employment? We do not find any difference in overall employment growth between country-industries that adopted robots at a high rate and those that lagged in robot adoption. How could this be, given that industrial robots necessarily replace human labour? We believe the answer is demand. Industries that adopt robots experience an increase in productivity, and hence a decrease in costs, which means they will sell their products at lower prices¹². Buyers of these products will respond by demanding larger quantities. To supply the larger quantities, more human workers need to be hired, since not all tasks can be done by robots. In this way, overall employment may not change (at least relative to non-adopting industries)¹³.

While overall employment, at the level of country-industry cells, appeared to be unaffected, there was clearly a disruptive effect on human labour from robot adoption, in that workers had to be moved to different tasks. This could be done by reassigning existing workers, and by firing some and hiring new ones. We find evidence for firing and hiring, as robots changed the composition of employment. In particular, employment of high skill workers (those with a college degree) grew faster in country-industries that adopted robots at a higher rate, while the opposite was true for low-skill workers. The same finding holds for the distribution of the wage bill across these workers. Thus, a sub-group of workers did experience an adverse impact from robot adoption.

The impact of industrial robots on productivity was positive and substantial. We conservatively calculate that on average, the increased

12. We indeed find a negative effect of robots on output prices.

13. Of course, manufacturing as a whole has seen employment declines in all developed countries. Our results speak to the comparison of country-industries with high robot adoption to those with low adoption. The shift of employment from manufacturing to services may also ultimately be due to technological change, but there is an important interaction with consumer preferences (Herrendorf et al. 2014).

use of robots contributed about 0.37 percentage points to the annual growth of GDP, which accounts for more than one tenth of total GDP growth over this period. The contribution to labour productivity growth was about 0.36 percentage points, accounting for one sixth of labour-productivity growth. This makes robots' contribution to the aggregate economy roughly on par with previous important technologies, such as the railroads in the nineteenth century (Crafts 2004) and the US highways in the twentieth century (Fernald 1999). The effects are also fairly comparable to the recent contributions of ICT (see e.g. O'Mahoney and Timmer 2009)¹⁴.

Robots are currently confined mainly to the manufacturing sector, although they are increasingly used in agriculture and construction, too. If robot capabilities continue to improve and expand, then robots may spread also to the services sector, and will thus contribute to productivity growth in a much broader way. This suggests that some of the recent concerns about a possible slowdown in productivity growth (e.g. Gordon 2016) may be overly pessimistic. Our findings do, however, come with a note of caution: there is some evidence of diminishing marginal returns to robot use, so robots are not a panacea for growth. Moreover, the rise of robots may not be positive for all, given that less-skilled workers may lose out.

Is this time different?

For most of history, technological change appears to have been a blessing overall: while it tends to disrupt labor markets by making some jobs obsolete, workers usually manage to adjust and find new jobs, and living standards rise. However, breakthroughs such as the driverless car or the supercomputer Watson (that beat the most able human contestants at the game Jeopardy!), which might have been unthinkable about fifteen years ago, prompt some observers to proclaim that “this time is different”. That is, technological change has now reached such a fast

14. Each of these technologies contributed more to productivity growth than did robots. But robots make up just over 2% of total capital, which is much less than previous technological drivers of growth. Relative to their share in total capital, the contribution of robots appears very similar to those of earlier innovations.

rate that the economy will not keep up with creating new jobs, and technological unemployment on a large scale will finally materialise (Ford 2016). In a more measured assessment, Autor (2015) acknowledges that innovations such as machine learning go some way to overcoming Polanyi's paradox (because machine learning algorithms do not rely as much on explicit instructions as conventional software does), but that progress and implementation will be slow, and labour markets will manage to adjust. Clearly there are plausible arguments supporting either view, so it seems a daunting task to predict with any confidence what human work will look like in say, 30-50 years from now.

As persuasively argued by Autor (2015), a general pattern that is very likely to continue is that technology substitutes for human labour in some tasks but not in others, and that new technologies complement workers in many of the remaining tasks. Perhaps the two key questions are, at what pace will the set of automatable tasks expand?, and what kinds of worker will be affected? While the first question is very hard to answer with any degree of certainty, there are some emerging trends that go some way in addressing the second question. In particular, it appears that high-skill workers are increasingly affected by automation. For instance, text analysis software now aids lawyers in pre-trial research, and surgical robots are employed in complicated procedures. In each case, the involvement of a highly skilled human lawyer or surgeon is still essential. But, as technology allows these workers to complete more assignments in a given amount of time, overall employment may fall, unless, as in the case of the industrial robots, demand responds sufficiently strongly to the lower price of the service. At the opposite end of the skill spectrum, tasks that have traditionally been the domain of untrained workers, such as driving or cleaning, also appear to be increasingly susceptible to automation. However, the mere feasibility of automating such tasks will not be enough to bring about the demise of the occupations concerned, given the prices of driverless cars and cleaning robots must also fall sufficiently relative to the cost of untrained human labor, which can be very low.

What are the implications for policy?

What, if anything, should public policy do to alleviate any disruptive impacts of technological change on labour markets? When someone chooses to train for a particular occupation, they take on the risk that automation could make this occupation obsolete before they reach retirement age. If, as appears to be the case, markets fail in supplying insurance for this type of risk, then such insurance should be provided by governments, for instance in the form of re-training subsidies. Furthermore, governments should ensure that the education system is responsive to the changing demands placed on workers. Finally, there may be a need for increased redistribution if the overall gains from technological change are distributed very unequally. These policy recommendations do not depend on the precise rate or direction of technological change, which are very difficult to predict. Rather, they are meant to address what appears to have been a constant feature of technological progress: a disruptive impact on the labour market combined with gains in wealth and living standards which sometimes are unevenly distributed¹⁵.

A different challenge for policy comes from the increased prevalence of ‘alternative work arrangements’, in particular independent contracting as facilitated by modern communication technologies, a phenomenon sometimes dubbed ‘the rise of the sharing (or gig) economy’¹⁶. It may no longer be appropriate to tie the social safety net as closely to traditional employment relationships as is often the case in developed countries. However, as long as jobs – whether of ‘traditional’ or ‘alternative’ forms – are available to those who wish to work, there is no reason to break the link between benefits and contributions, as would be the case under a basic income guarantee.

Perhaps the biggest challenge in the years to come is not posed by technological progress itself, nor, say, globalisation and mass immigration – but by a political climate that is not conducive to constructive

15. Incidentally, the same policies would also address some of the challenges posed by greater openness to trade and globalisation.

16. Katz & Krueger (2016) document a rise in the fraction of workers engaged in alternative work arrangements from 10% in 2005 to 16% in 2015 in the US.

policy discussions. Two developments in particular are worrying. First, there has been a backlash against free trade and immigration – witness the recent surge of populist parties and movements across developed countries – despite much evidence demonstrating that the gains of trade and immigration outweigh the losses¹⁷. Should we then expect a similar backlash against technological progress, even if it continues to largely be a blessing? Second, many popular ‘remedies’ being debated tend not to focus on the actual problem at hand – for instance, restrictions on trade seem to be favoured by many over targeted help for affected workers, minimum wages over income tax credits to help low-wage earners¹⁸, or a basic income guarantee over a social safety net that is better tailored to deal with the risks of technological change and globalisation. There is no shortage of promising policy proposals. However, a better understanding of the current political climate is needed to develop strategies that ensure that these proposals receive more attention than their inferior counterparts.

17. A prime example is the UK Brexit vote, despite research showing substantial benefits of EU trade integration (Dhingra et al. 2016) and immigration (Wadsworth 2015).

18. See Burkhauser (2015).

References

Acemoglu, Daron, 2002, 'Technical Change, Inequality, and the Labor Market,' *Journal of Economic Literature*, vol. 40(1).

Acemoglu, Daron & David Autor, 2011, 'Skills, Tasks and Technologies: Implications for Employment and Earnings,' Chapter 12 in *Handbook of Labor Economics*, volume 4b.

Autor, David, 2013, 'The "task approach" to labor markets: an overview,' *Journal of Labour Market Research*.

Autor, David, 2015, 'Why Are There Still So Many Jobs? The History and Future of Workplace Automation,' *Journal of Economic Perspectives*, 29(3).

Autor, David, Lawrence Katz & Melissa Kearney, 2006, 'The Polarization of the U.S. Labor Market,' *American Economic Review*, 96(2).

Autor, David, Frank Levy & Richard Murnane, 2003, 'The skill content of recent technological change: an empirical exploration,' *Quarterly Journal of Economics*, 116(4).

Breman, Anna, 2016, 'Diginomics and the Productivity Puzzle', this volume.

Brynjolfsson, Erik & Andrew McAfee, 2014, *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, W.W. Norton & Company.

Burkhauser, Richard, 2015, 'The minimum wage versus the earned income tax credit for reducing poverty,' *IZA World of Labor*.

Coase, R. H. 1937. 'The nature of the firm', *Economica*, 4(16): 386-405.

Crafts, Nicholas, 2004, 'Steam as a General Purpose Technology: A Growth Accounting Perspective,' *The Economic Journal*, 114(495).

Dhingra, Swati, Gianmarco Ottaviano, Thomas Sampson & John Van Reenen, 2016, *Brexit: the impact on UK trade and living standards*, Centre for Economic Performance.

Fernald, John, 1999, 'Roads to Prosperity? Assessing the Link between Public Capital and Productivity,' *American Economic Review*, 89(3).

Ford, Martin, 2016, *Rise of the Robots: Technology and the Threat of a Jobless Future*, Basic Books.

Goldin, Claudia & Lawrence Katz, 2008, *The Race between Education and Technology*, Belknap Press.

Gordon, Robert, 2016, *The Rise and Fall of American Growth*, Princeton University Press.

Goos, Maarten & Alan Manning, 2007, 'Lousy and Lovely Jobs: The Rising Polarization of Work in Britain,' *The Review of Economics and Statistics*, 89(1).

Goos, Maarten, Alan Manning & Anna Salomons, 2009, 'Job Polarization in Europe,' *American Economic Review*, 99(2).

Graetz, Georg & Guy Michaels, 2015, *Robots at Work*, CEPR Discussion Paper 10477.

Herrendorf, Berthold, Richard Rogerson & Akos Valentinyi, 2014, 'Growth and Structural Transformation,' chapter 6 in *Handbook of Economic Growth*, vol. 2B, Elsevier.

Katz, Lawrence & Alan Krueger, 2016, *The Rise and Nature of Alternative Work Arrangements in the United States, 1995-2015*, working paper.

O'Mahony, Mary & Marcel Timmer, 2009, 'Output, Input and Productivity Measures at the Industry Level: The EU KLEMS Database,' *Economic Journal*, 119 (538).

Nordhaus, William, 2007, 'Two Centuries of Productivity Growth in Computing,' *The Journal of Economic History*, 67(1).

Polanyi, Michael 1966, *The Tacit Dimension*, New York: Doubleday.

Wadsworth, Jonathan, 2015, *Immigration and the UK Labour Market*, CEP Election Analysis Papers 019, Centre for Economic Performance, LSE.

Darja Isaksson is an entrepreneur and advisor within consumer experience design, innovation strategy and data analytics. She was awarded ‘Most influential opinion maker of the year’ in 2016 by Swedish business magazine *Veckans Affärer*, and is currently a member of the Swedish Prime Minister’s innovation council. **Karl Wennberg** is Professor in Business Administration at the Stockholm School of Economics, deputy director and co-founder of the Institute of Analytical Sociology (IAS) at Linköping University. IAS employs big data to perform detailed empirical research on complex social, political, and cultural processes.

Digitalisation and collective value creation

This chapter discusses the spread and impact of digitalisation as a disruptive technological change. We show how digitalisation is intimately connected to globalisation by first, being dependent on globalisation for its impact, and second, enhancing the speed of globalisation. Digitalisation lowers barriers to funding, marketing, sales and distribution, and enables an increasing global flow of goods, services, and financial transactions. We discuss how digitalisation also contributes to changing consumer habits and a blurring line between producers and consumers, where the latter now have capabilities to build collective knowledge by themselves becoming producers. Digital platforms are emerging, aggregating data and providing new business models where contact costs are approaching zero. These platforms wield

strong economic power and the algorithms by which they operate also change incentives and transaction costs for producers and consumers. We sketch the patterns by which industries digitalise as being characterised by one or a few ‘platforms’ dominating a global market, but where such platforms also facilitate the emergence of narrower niche businesses and products and allow new types of micro-multinationals to reach out to a larger global crowd and satisfy latent demand. These changes have already happened in media and music, and the principles seen in these industries can be seen as emerging in other sectors. We conclude by highlighting the potential of digitalisation to enhance the value of collective goods. We particularly highlight the cases of health care and energy, and discuss how digital technologies can contribute to collective value creation in these areas.

Digitalisation and globalisation

Digitalisation can be defined as the use of digital technologies to change a business model and provide new revenue and value-producing opportunities. It’s also the process of moving to a digital business. Digitalisation is a technological force that enhances globalisation in both economic and cultural ways. The effect of digitalisation on increased globalisation can be seen in a number of areas, including digital goods, enhanced cross-border communication and globally distributed teams, electronic platforms, and optimised flows of goods and services. Together, such changes allow individuals and organisations to share information and knowledge more rapidly and seamlessly than any time before in human history, facilitating the growth of collective knowledge¹.

The effects of digitalisation on economic and cultural globalisation are however not uncontroversial, nor unidirectional. Given slower productivity growth, some people have been debating the extent to which digitalisation is hype or actually a real force². Further, the rela-

1. While earlier technological innovations such as written language, the printing press, or the telegraph has also magnified the spread of knowledge, none of these have diffused as rapidly, nor allowing individuals to interact collectively on the same scale as digitally enabled technologies.

2. Wolf, M. 2014. ‘Same as It Ever Was: Why the Techno-optimists Are Wrong’. *Foreign Affairs*. 94(4). July/August.

tionship between digitalisation and globalisation is not necessarily unidirectional, but may grow in tandem.

Globalisation was underway before and during the rise of the internet, enhanced by increased travel and cultural exchange, in addition to political changes, such the fall of the Soviet Union, China's rise as an economic power, and several important regional and global free-trade agreements. In several respects, the long-term globalisation following World War II has provided economic, regulatory, and cultural support for the digitalisation of goods and services. For example, we now have a global language – English. Student exchanges are frequent at high school and university levels alike. Expatriate workers and migrant workers are more common than ever: more than 250 million people today reside outside their country of birth. The EU's internal labour market has facilitated cultural exchange and is to date the most popular among the EU pillars³. 44 million are engaged in cross-border online work using digital tools, and over 360 million participate in cross-border e-commerce⁴. But, economic exchange is but one of the tokens by which globalisation and digitalisation are interconnected. Almost one billion people around the world have one or several international social media connections, and those connections are being put to use in various ways⁵. As a society, people, places and things are now becoming interconnected in a way that is radically changing the way we operate anything from trade and transport, to education and health care.

Globalisation of digital goods

Measuring digital trade and its impact on globalisation is complex⁶. However, the flow of cross-border data stemming from trade in digital goods constitutes a readily definable and easy-to-understand portion

3. Barslund, M. & Busse, M. 2014. 'Making the Most of EU Labour Mobility'. Centre for European Policy Studies, Brussels, ISBN 978-94-6138-407-2.

4. McKinsey & Co. 2016. 'Digital globalization: The new era of global flows'. March 2016: McKinsey Global Institute.

5. Ibid.

6. OECD (2013). 'Measuring the Internet Economy: A Contribution to the Research Agenda', *OECD Digital Economy Papers*, No. 226, OECD Publishing. <http://dx.doi.org/10.1787/5k43gig6r8jf-en>.

of growing international digital exchange. Music, books, videos, games and advertising are examples of industries where goods are often distributed digitally. In such industries, digitalisation has radically impacted business models as well as market structure. Brick-and-mortar chains such as Blockbuster and Borders have been replaced by Amazon, HBO, Netflix, and Spotify. Once digital disruptors, they are now major players with global reach, and often rapid growth. However, the same global players both enable and inspire new competition as well as business aimed at globally distributed niche consumers. For instance, automated advertisement services such as Sweden's Widespace allows companies to tailor advertisement directly to consumers and track consumer engagement. Splay is a network of Youtube stars, providing advertisement opportunities to brands, and entertainment to consumers. We see that digital goods simultaneously facilitate certain 'mono-cultures' (homogenisation of mainstream demand) along with the proliferation of a host of global niches and micro-communities (heterogeneous demand).

Globally distributed electronic platforms

Digitalisation also impacts new startups, when digitally enabled education, crowdfunding, e-commerce, and global access to talent are combined. Online platforms such as Kickstarter, Fundedbyme and Kivra provide funding and marketing for a large pool of small businesses, varying from traditional trade of goods and services to high-tech innovators all over the world. Social media platforms enable global reach in marketing for even small companies, enabling them to become micro-multinationals. E-commerce platforms such as Amazon, Avito and Alibaba, digital payments through Paypal, Klarna and Stripe and sharing-platforms like Airbnb or Elance, provide individuals and entrepreneurs with a low-cost solution and the potential for global reach for their goods and services. The growth of globally accessible electronic platforms decreases the cost of international interactions and transactions, which in turn facilitates the creation of new markets and user

communities of global scale⁷. For instance, 90% of commercial sellers on eBay export to other countries, compared to less than 25% of traditional small businesses⁸. Further, businesses selling through eBay are also more likely to export to many countries, and less likely to lose their export-destination relationships than traditional businesses⁹. Globally accessible electronic platforms are enhancing the globalisation of people, transactions and businesses everywhere.

Globalised talent

Digitalisation is also changing the flows of skills and people. One way this digitalisation is doing that is by enabling remote work across borders. For large multinational companies, costs and security concerns have hitherto been important factors driving the growth of globally distributed teams that work together remotely across borders. But access to, and cost of, talent pushes startups as well as new NGOs and social movements to bring together talent from all over the world. Distributed teams are now often a default mode of organisation in many sectors: these include multinational corporations, digital service providers and resource-constrained startups¹⁰. Distributed teams can be cheaply organised using digital tools such as Dropbox, Box and Slack for remote collaboration, or Intercom and Zendesk for customer-relationship management¹¹. This development parallels the overarching trend of more and more people working as freelancers or working for an employer from their home. McKinsey & Co reports that 58% of U.S. companies expect to use more tempo-

7. McKinsey & Co. 2016. 'Digital globalization: The new era of global flows'. March 2016: McKinsey Global Institute.

8. Ibid.

9. Lendle, A., M. Olarrega, S. Schropp, P.-L. Vézina. 2013. 'eBay's anatomy'. *Economics Letters*, 121(1): 115-120.

10. Oshri, I., Van Fenema, P., & Kotlarsky, J. 2008. 'Knowledge transfer in globally distributed teams: the role of transactive memory'. *Information Systems Journal*, 18(6): 593-616.

11. It should be noted that as the cost of communicating over long distances is decreasing, the marginal value of physical interaction is *increasing*, so that the most productive firms and individuals increasingly chose to co-locate in clusters. Moretti, E., 2012. *The new geography of jobs*. Houghton Mifflin Harcourt.

rary arrangements at all levels, production, clerical, and managerial¹². Freelancers work onsite, from their homes, and through co-working spaces such as the US WeWork, which fill the human need for social interaction at work. These developments are amplified by digital services such as Krop, Guru and 99designs, which reduce the cost of contact and make it easy for individuals to find projects as well as contractors. Companies and organisations can now access a global pool of talent working in areas varying from programming and design to management, finance and legal services. Not only does this provide opportunities for increased cost efficiency, but it's perhaps primarily a driver for enabling economies of latent demand¹³.

Digitally improved products

Digitalisation has a significant impact on physical goods and complex supply chains. Containers, trucks, cars, home appliances and other machinery today comes with embedded sensors, enabling optimisation of flows from production to supply chain and service. This is changing traditional supply chains. For example, digitalisation changes supply chains by providing real-time data of the flow of goods across different modes of transportation, as well as increasing automation of these modes¹⁴. Logistics companies can already track and coordinate their ships, trucks, and containers across borders, letting customers track their parcels online. General Electric (GE) is an example of using this to improve their business, as they've changed one of their business models from selling airplane engines and later charging for the maintenance of them, to charging their customers for the value of optimising air fleets. GE has also launched a platform called Predix, where other products and services can upload data, and a world of developers can develop services based on that data. Similar development can be seen in areas varying from energy to agriculture and health care. The use of sensors, big data analysis and artificial in-

12. Miller, J. G., & Miller, M. 2012. 'The rise of the supertemp'. *Harvard Business Review*, 51: 4-12.

13. Susskind & Susskind, 2015. *Future of Professions*.

14. *SCDigest Supply Chain Digitization Benchmark Survey 2016*.

telligence now allows complex networks of production and consumption to increase visibility and begin to be optimised on a scale that was not possible until technology development just recently made the technology cost efficient enough. As a result, entire business ecosystems are now changing; increasing in complexity, but also lowering the barriers for new actors to add value through data, services and new products.

Globally accessible collective knowledge

Wikipedia was one of the world's first really successful examples of knowledge-generating open platforms, but knowledge is now increasingly being made available to a global audience through both closed and open knowledge platforms as well as open source solutions for the creation of new goods and services. Open knowledge-platforms such as Patientslikeme allow people to share information and experiences on personal health as well as advanced treatment. The value of the data and people collected by Patientslikeme is acknowledged by companies such as AstraZeneca, now partnering with the platform for clinical studies. Github, an open source code repository, now has 14 millions developers sharing code and knowledge, and it's currently the largest host of source code in the world. More and more companies are using similar principles of open collaboration as a part of their business model. For example, Sweden's Propellerhead Software – a niche producer of electronic music production tools – rapidly grew to number two worldwide and was able to gather a large customer base which started to produce new software and launch businesses through the open-source solutions provided by Propellerhead¹⁵. Producers are co-creating new products, and consumers are becoming producers.

Non-profit and educational institutions are also contributing to globally accessible collective knowledge. Millions of people

15. Autio, E., L. Dahlander, L. Frederiksen. 2013. 'Information Exposure, Opportunity Evaluation, and Entrepreneurial Action: An Investigation of an Online User Community'. *Academy of Management Journal* 56(5) 1348-1371.

are regularly attending online courses – often in the form of Massive Open Online Courses (MOOCs) – as a part of their higher education, or part of life-long learning¹⁶. Massachusetts Institute of Technology (MIT) has opened the content of much of their traditional courses for the general public, making materials used in the teaching of MIT's subjects available on the Web. This allows anyone anywhere to access online lectures, course material, etc. Now, Coursera, HarvardX and MIT are all also experimenting with various ways of charging for awarding credits for such courses. Many smaller colleges and universities are taking advantage of such innovations to enhance their educational offering, and transition economies such as India are linking up to MOOCs and integrating this to scale up their institute of higher learning to large populations that previously had no access to them. The types of knowledge exchange solutions are part of the 'sharing economy' where resources are distributed and shared, with value appropriability non-governed¹⁷. On an individual level, pretty much everybody regularly use open knowledge without thinking much about it. A recent survey by Google showed that 91% of smartphone users turn to their smartphone for ideas while doing a given task¹⁸, which when seen as a mass behaviour is a relevant phenomenon in understanding the potential of digitalisation for the global diffusion of knowledge. As a result of digitalisation, more people than ever before are now able to participate both in the creation and distribution of knowledge, and costs for finding and co-creating complex solutions to existing problems are decreasing¹⁹.

16. Stadin, E. 2016. 'Öppna nätkurser för tusentals: Utvecklingen av moocar i Europa' FORES Policy paper 2016:1.

17. Felländer, A., Ingram, C. & Teigland, R. 2015. 'Sharing Economy; Embracing Change with Caution'. Entreprenörskapsforum.

18. Google/Ipsos, Consumers in the Micro-Moment study, March 2015. Based on the online population n=9598.

19. Ballantyne, P. 2014. *Challenge Prizes: A practice guide*, NESTA.

Platform-based ecosystems

Accepting the notion that digitalisation is reshaping almost all industries, common patterns can be seen in transformations that have already happened, and those that are beginning to emerge. It's clear that platforms with global impact lower transaction cost and barriers to financial investments, marketing and distribution of digital goods, services and knowledge²⁰. Algorithms optimise and manage flows, and the patterns that algorithms take into account are often based on consumer behaviour data. Further, digitalisation facilitates collective value creation in many sectors. The consequences have reshaped the structure of affected industries, their competitive conditions, and consumer experiences.

The first industries to display this pattern were news and entertainment, with digitalisation making news and content more readily available to consumers. The cost of distribution plummeted and consumers today have access to more news, music and content than ever before. In parallel, a handful of platform actors have become pivotal to anyone seeking to reach out or build a business in these sectors. Social media platforms and search engines now dictate the terms for most of that industry, and that their algorithms have profound impact on consumption patterns. In this digital ecosystem, geographical boundaries no longer dictate competitive conditions in an industry, while the reach and relevance of individuals and organisations to a globally widespread audience do. This has had a profound impact on providers of digital content everywhere. *The Wall Street Journal* or *Vogue* are two examples of content providers that have benefited from a growing audience of connected English-speaking consumers, quite likely at the expense of more local content providers, such as national papers in non-English speaking regions. This is in itself an interesting phenomenon, since it's likely to add to the globalisation of ideas and behavior in a connected population. At the same time, we have seen numerous examples of how small actors have used these platforms to gain access to large audiences, and in turn, build their own business. Often in niches

20. Choudary, S.P., Van Alstyme, M.W., Parker, G.G., 2016. Platform Revolution: How Networked Markets Are Transforming the Economy—And How to Make Them Work for You. New York: WW Norton & Company.

that previously didn't exist, or could not easily be capitalised upon.

We've seen a similar development in the mobile phone industry. The biggest innovation of Apple related to the iPhone was the App Store. It allowed masses of developers to contribute value to Apple's customers by building mobile apps. Consumers benefited immensely from this new rich ecosystem of mobile services that the app store enabled, and it has radically changed behaviour and businesses all over the world. This innovation was key to helping Apple achieve a remarkable success; Apple pushed an entire industry from a situation where five major mobile phone manufacturers (Nokia, Sony Ericsson, Motorola, Samsung and LG) collectively controlled 90% of the industry's profits, to a situation where Apple dominated completely. In 2015, the iPhone alone represented 92% of the industry's profits²¹.

What characterises companies that have successfully achieved a platform position, is that the unique resource they have access to is massive amounts of platform participants; both producers and consumers. And they manage to do so, and not necessarily by having access to unique resources or production systems.

In sum, digitalisation so far has led to rapid oligopolisation of industries across borders where large platforms such as Apple's App Store, Google and Facebook dictate competitive conditions²². But at the same time, digitalisation has also facilitated the emergence of more narrow 'niche' markets, where formerly hobbyists or unprofitable companies can reach out to a larger global crowd and satisfy formerly latent demand²³.

These types of digital-induced changes are now beginning to show up in industries, such as transportation, energy and health. Emerging digital platforms indicate a change in the nature of competition²⁴ due to lower barriers and access to a global audience, combined with the use of

21. Forbes, 2016. 'Apple's iPhone Profits Will Weed Out Other Players'.

22. Power, D. & Hallenruetz, D. (2007). 'Competitiveness, local production systems, and global commodity chains in the music industry: Entering the U.S. market'. *Regional Studies*, 41(3), 377-389.

23. Lopes, P. D. (1992). Innovation and diversity in the popular music industry, 1969-1990. *American Sociological Review*, 57(1), 56-71., and Ahrens, S. & Kreidenweiss, A. (2012). *Industry-wide business model innovation: The case of the Swedish music industry*. (Master's Thesis, Stockholm School of Economics).

24. Van Alstyne, Parker & Choudary (HBR April 2016), 'Pipelines, Platforms and the New Rules of Strategy'.

consumer behaviour data. These sectors are crucial for the transformation to a more sustainable society, since they are sectors which infrastructure, and potentially also production and distribution, are characterised as public goods. In the last part of this chapter, we therefore dwell on the cases of transportation, energy and health, and what digitalisation could mean for higher value creation in public goods.

Transportation

Transportation in the industrialised world is wasteful. An average city today devotes about 50% of its space to roads and parking facilities, yet the average European car is used for driving only 5% of the time²⁵. At the same time, transport constitutes 30% of our fossil-fuel emissions. Transportation of goods as well as people is pivotal to achieve a vision of a fossil-free society. Digitalisation is needed to address both the changes in behaviour that's needed, and to optimise use of available resources.

Within the transport sector, new actors are leading the way and targeting platform positions in a new ecosystem of transportation services. Companies such as Waze and Uber demonstrate the transformative power of digitally distributed services, quickly reaching a global audience. As an enabler of the 'gig economy', Uber is often used as the most typical example of the new breed of platform businesses and their impact on traditional markets²⁶. But as part of a growing number of transport services, Uber is also one of the actors that create alternatives to the growing number of city people who value convenience of usage above convenience of ownership. Waze, a collaborative tool for route planning and collective information-sharing is now used by as much as 10-15% of the population in large, densely populated cities such as Los Angeles and New York, and the service is impacting transportation patterns there. At the same time, Tesla is challenging the car industry by building electric vehicles, but also by building cars that are contin-

25. Ellen McArthur Foundation, 2015. *Growth within; a circular economy vision for a competitive Europe*.

26. Laurell, C. & Sandström, C. 2016. 'Analysing uber in social media - Disruptive technology or institutional disruption?' *International Journal of Innovation Management*, 20(5): DOI: <http://dx.doi.org/10.1142/S1363919616400132>.

uously updated, meaning their customers purchase a car but a digital service related not only to maintenance but even advanced functionality such as autonomous driving algorithms and continuously updated security functions. Previous examples from the telecom industry show how quickly such experiences impact consumer expectations on what a product should deliver.

However, while single actors provide relevant innovations, a sustainable transport system will require various means of transport to be integrated into services for consumers. In order to enable this, nations such as the Netherlands and the UK are currently using digitalisation to innovate its transport sector by providing open national platforms for data and digital services²⁷. Such platforms enable new actors to build services that make travelling easier by creating route plans that integrate various forms of transport.

A transport system that is fully integrated in that way also highlight the need for a new type of actor; one that has a mandate over algorithms that steer incentives in choice of route as well as means of transport. Dynamic pricing could be applied to trips, taking into account supply and demand for parking as well as road accessibility. An interesting established solution of road accessibility is digital startup Waze, now used by up to 10% of all drivers in major US cities. By showing how traffic actually flows in real time, Waze is also generating data used by public authorities when planning for new infrastructure. But private startups are not the only actor of relevance. For example, Dutch authorities are aiming for another solution; one where transport data is treated as part of public goods.

Energy

With renewable energy becoming increasingly cost effective and storage solutions beginning to catch up, consensus around the basics of tomorrow's energy systems is starting to emerge. Energy systems are expected to become an increasingly complex mix of small- and large-scale

²⁷. www.plannerstack.org.

producers, where distributed systems of energy production, as well as energy storage, become predominant standards²⁸. Large groups of consumers will be ‘prosumers’, partly passively through using products where smart-energy harvesting solutions are an inherent part of the product, but also actively by providing and storing renewable energy²⁹.

Actors such as General Electric (GE) with the Predix platform and Tesla’s electronic cars now combine electric vehicles with solar energy and storage appear to be aiming for platform positions in a new ecosystem, where energy and transportation are intimately connected, and where small-scale producers and storage solutions are key to an overall system of optimised production and access. In the Netherlands, Germany and the US, hospitals and high-tech companies such as Apple and Facebook, are increasing investment in becoming self-sufficient in energy production, and Apple is even beginning to sell its surplus. At the same time, companies such as Sonnenbatterie in Germany and Open Utility in the UK are enabling small-scale producers and consumers to trade energy directly with each other, and in some places consumers are simply taking things into their own hands³⁰.

The energy sector is highly regulated, and countries have chosen different paths as to how much this development should be encouraged. Regardless of that, it’s clear that the cost of renewable, as well as energy storage, is dropping, and that it has an impact on consumer expectations and behaviour. In a scenario where companies with technology aimed directly at consumers, such as Tesla and Sonnenbatterie, have a big impact, it’s clear that companies in markets where energy prosumers are empowered early have a head start. Perhaps there are parallels to previous experiences of what enabling a population to become early adopters of internet and social media has meant in terms of innovative companies, as well as creating a mature and relevant market for companies that push ahead.

28. Luo, X., Wang, J., Dooner, M., & Clarke, J. 2015. ‘Overview of current development in electrical energy storage technologies and the application potential in power system operation’. *Applied Energy*, 137: 511–536. SETIS. 2014. *Set plan: Energy Integrated Roadmap*. SETIS (Strategic Energy Technologies Information System). European Commission. <https://setis.ec.europa.eu>.

29. Gerhardt, W. 2008. *Prosumers: A New Growth Opportunity*. Cisco Internet Business Solutions Group (IBSG). <http://www.cisco.com/go/ibsg>.

30. <http://www.npr.org/sections/alltechconsidered/2016/07/04/482958497/how-blockchain-helps-brooklyn-dwellers-use-neighbors-solar-energy>.

Health

The health sector is another sector that's crucial for achieving sustainable societies and where digital transformation is beginning to change prevailing practices and solutions. The need for increased quality and efficiency in health care is an increasing challenge for all countries, be it because of increasing costs due to longer lifetimes and increasingly complex health conditions, or because of a growing need for quality health care in rural areas. There is significant value about to be created by those who manage to provide scalable, efficient solutions. Any digital solutions will require data, and digital platforms for distribution. Data aggregators, such as patientslikeme.com, Apple's research kit, and the GE Predix platform already make valuable health data available to researchers and innovators. Digitally distributed health services, as well as medtech innovations, are growing. These vary from tools for remote x-ray analysis to early diagnosis of Alzheimer's and regular primary health care services, such as kry.se and mindoktor.se in Sweden³¹. In Singapore, most citizens regularly schedule doctors' appointments and handle their prescriptions through the social media platform Wechat. In China, policymakers envision using digitally distributed health services as a way to increase availability for large groups in society. At the same time, modern travel habits have helped increase a global trend towards health related tourism, varying from dental care to specialised cancer treatment and surgery. Nationally, this development is slowed down by issues related to technical legacy, integration cost and sometimes outdated regulations. The issue of data interoperability between systems and continuous development of regulation is therefore of high priority in most countries. This is an area where Sweden is currently lagging behind forerunners such as Denmark and Singapore. But there's also a strong international aspect to this, and with legal and institutional constraints currently hampering the flow of data across borders, digitally-enabled health services may benefit even the more if regulatory frameworks were updated.

31. Black, A. D., Car, J., Pagliari, C., Anandan, C., Cresswell, K., Bokun, T., McKinstry, B., Procter, R., Majeed, A., & Sheikh, A. 2011. *The impact of eHealth on the quality and safety of health care: a systematic overview*. *PLoS Med*, 8(1): e1000387.

Data produced by individuals is becoming more valuable to both researchers and consumer-oriented innovative companies, and patterns seen through big data are key to optimising both flows in health care and clinical studies. We are still in the early stages, but large amounts of data is starting to flow between various digital services and between actors in both the public and private sector. Current examples highlight the increasing potential to offer health-related services as digitally distributed services, as well as blended services where the cost and effort for a physical meeting make it worthwhile to travel. Our forecast is that the increasing specialisation among advanced treatments, together with globalisation of knowledgeable consumers, means digitalisation of health care will lead to a growing global market for high-quality and specialised health services, as well as digitally distributed services across borders for less-complex ailments and services, all in a way that maximise the value of both online and physical services. But, most importantly, countries that now act in a way that enables the broad landscape of international innovators to benefit their health care system by providing solutions to it while maintaining quality control, will see considerable benefits in both improved efficiency and quality of health. Such markets are also more likely to create the winners in tomorrow's ecosystem of digital health systems. However, it's unclear if countries will prevent the emergence of globally dominant players with power over both data, platforms and algorithms, such as has been seen in the media and telecom industry.

Security and integrity

Since the impacts of digitalisation mentioned above are based on services using data, issues related to security and integrity are crucial and worth highlighting, though in-depth discussion of these topics is out of the scope of this chapter. While already aware of the dangers of misuse of data, most people today share information about their behaviour and opinions and location in exchange for increased relevance and more-efficient service in areas varying from route planning to advertising. This can be seen as naïve and dangerous. However, taking into ac-

count the growing amount of options and messages that already bombard individuals using digital devices, using data in exchange for better relevance and less friction can also be seen as rational. It's clear that digitalisation is raising issues concerning security and integrity for individuals, as well as corporations. The rapidly expanding digitalisation of health and transportation highlights that in exchange for personal data, individuals may also benefit from collective goods, such as environmental-friendly energy transportation, decreased traffic congestion, and more accessible health care services. A key challenge of our time is therefore to find ways to enable individuals, organisations and countries to reap the benefits that digitalisation can give, while doing so in a secure manner that also protects society and individuals from the potential dangers of massively accessible data. 'Integrity' is a concept that has meaning in a cultural context. As a society, we need to find out what a functional concept of integrity is in a world full of, and fueled by, data. This is no easy task, and is being addressed by policy makers as well as companies aiming to find a good balance between value and risk.

Conclusions – leveraging collective goods

The examples of transport, energy and health care highlight the potential of digitalisation to redefine and enhance the production and consumption of collective goods. Collective goods are those judged as having 'collective value', in that individuals cannot be effectively excluded from use and where use by one individual does not reduce availability to others, such as the natural environment, central public infrastructure, and in practice also basic education and healthcare. Our chapter has focused specifically on the potential of digital technologies to enhance value creation, sharing, and increased productivity in health care and among environmental-friendly technologies. Examples from other industries of how digitalisation may enhance the production and consumption of collective goods could be found or envisioned. While digitalisation offers the potential for radically higher value creation in the production and distribution of public goods, measuring the impact of digitally enabled public goods remains in its infancy since productiv-

ity could be underestimated due to macroeconomic models not being able to gauge all new types of value creation³².

A central question for legislators and policymakers are the extent to whether these changes are all for good, and for whom? The emerging global ecosystem standard as seen in computers, software, mobile handsets and cars highlights that speed-to-market and economies of scale may lead to a limited number of global actors, such as Google handling services that were previously handled by nationally or regionally separated actors. In the cases of energy or health care, one can easily expect regional and national utility companies or health care providers to be replaced by multinational companies. Legislators in EU and national governments must consider what steps can be taken that enable a large variety of companies and citizens in their respective constituencies to make the leap towards digitalisation, and not be left behind. Since many digital businesses are based on collective value creation through big data, legislators should consider how to maximise the delivery of collective goods for sustainable development; be it by themselves or through private companies operating by authorised standards of data transportation, maintenance, integrity, and openness. Research on the impact of open data sharing in relation to all these areas so far has been scarce, but a study comparing available reports concludes that the findings indicate that government data openness positively affects the formation of knowledge bases in a country and that the level of knowledge base positively affects a country's global competitiveness³³.

Furthermore, authorities would benefit from finding ways to enable innovative products to add value to public systems through platform logic rather than through complicated and often inefficient public procurement of 'standalone' systems. Today, public procurement in the EU and its member states operate in regional and organisational silos with expensive transaction costs and 'lock-in' effects³⁴.

The development we outline is still at its nascent stage and expected

32. Felländer, A., Fölster S. & Ingram, C. 2016. *Det datadrivna samhället*. Stockholm: Digitaliseringskommissionen.

33. Lee, J.-N., Ham, J., & Choi, B. 2016. *Effect of Government Data Openness on a Knowledge-based Economy*. *Procedia Computer Science*, 91: 158-167.

34. Osborne, S. P., & Brown, L. 2013. *Handbook of Innovation in Public Services*. Cheltenham: Edward Elgar.

to increase exponentially in the next few decades. However, it's clear that government initiatives in these areas are relevant now, and will have a big impact on innovation and public-sector quality and efficiency in the decade to come.

References

- Autio, E., L. Dahlander, L. Frederiksen.** 2013. 'Information Exposure, Opportunity Evaluation, and Entrepreneurial Action: An Investigation of an Online User Community'. *Academy of Management Journal* 56(5) 1348-1371.
- Ballantyne, P.** 2014. Challenge Prizes: A practice guide NESTA.
- Barslund, M. & Busse, M.** 2014. 'Making the Most of EU Labour Mobility'. Centre for European Policy Studies, Brussels, ISBN 978-94-6138-407-2.
- Black, A. D., Car, J., Pagliari, C., Anandan, C., Cresswell, K., Bokun, T., McKinstry, B., Procter, R., Majeed, A. & Sheikh, A.** 2011. *The impact of eHealth on the quality and safety of health care: a systematic overview*. PLoS Med, 8(1): e1000387.
- Choudary, S.P., Van Alstyne, M.W., Parker, G.G.,** 2016. *Platform Revolution: How Networked Markets Are Transforming the Economy – And How to Make Them Work for You*. New York: WW Norton & Company.
- Ellen McArthur Foundation,** 2015. Growth within; a circular economy vision for a competitive Europe.
- Felländer, A., Ingram, C. & Teigland, R.** 2015. 'Sharing Economy; Embracing Change with Caution'. Entreprenörskapsforum.
- Felländer, A., Fölster S. & Ingram, C.** 2016. *Det datadrivna samhället*. Stockholm: Digitaliseringskommissionen.
- Forbes,** 2016. Apple's iPhone Profits Will Weed Out Other Players. Online, 2016-02-13 (<http://www.forbes.com/sites/chuck-jones/2016/02/13/apples-iphone-profits-will-weed-out-other-players/#760acd4413e5>). Accessed 2016-10-25.
- Gerhardt, W.** 2008. Prosumers: A New Growth Opportunity. Cisco Internet Business Solutions Group (IBSG). <http://www.cisco.com/go/ibsg>.
- Lee, J.-N., Ham, J., & Choi, B.** 2016. Effect of Government Data Openness on a Knowledge-based Economy. *Procedia Computer Science*, 91: 158-167.

McKinsey & Co. 2016. 'Digital globalization: The new era of global flows'. March 2016: McKinsey Global Institute.

Miller, J. G. & Miller, M. 2012. 'The rise of the supertemp'. *Harvard Business Review*, 51: 4-12.

Lendle, A., M. Olarrega, S. Schropp & P.-L. Vézina. 2013. 'eBay's anatomy'. *Economics Letters*, 121(1): 115-120.

Laurell, C. & Sandström, C. 2016. 'Analysing uber in social media - Disruptive technology or institutional disruption?', *International Journal of Innovation Management*, 20(5): DOI: <http://dx.doi.org/10.1142/S1363919616400132>.

Lopes, P. D. 1992. 'Innovation and diversity in the popular music industry, 1969-1990'. *American Sociological Review*, 57(1), 56-71., and Ahrens, S. & Kreidenweiss, A. (2012). *Industry-wide business model innovation: The case of the Swedish music industry*. (Master's Thesis, Stockholm School of Economics).

Luo, X., Wang, J., Dooner, M. & Clarke, J. 2015. 'Overview of current development in electrical energy storage technologies and the application potential in power system operation'. *Applied Energy*, 137: 511-536.

OECD, 2013. 'Measuring the Internet Economy: A Contribution to the Research Agenda', OECD Digital Economy Papers, No. 226, OECD Publishing. <http://dx.doi.org/10.1787/5k43gg6r8jf-en>.

Osborne, S. P. & Brown, L. 2013. *Handbook of Innovation in Public Services*. Cheltenham: Edward Elgar.

Oshri, I., Van Fenema, P. & Kotlarsky, J. 2008. 'Knowledge transfer in globally distributed teams: the role of transactive memory'. *Information Systems Journal*, 18(6): 593-616.

Power, D., & Hallencruetz, D. 2007. 'Competitiveness, local production systems, and global commodity chains in the music industry: Entering the U.S. market'. *Regional Studies*, 41(3), 377-389.

SCDigest, 2015. 'Supply Chain Digitization Benchmark Survey' www.gartner.com/technology/supply-chain/top25.jsp. Accessed Sept. 25.

SETIS. 2014. Set plan: *Energy Integrated Roadmap*. SETIS (*Strategic Energy Technologies Information System*). European Commission. <https://setis.ec.europa.eu>.

Stadin, E. 2016. 'Öppna nätkurser för tusentals: Utvecklingen av moocar i Europa', FORES Policy paper 2016:1.

Wolf, M. 2014. 'Same as It Ever Was: Why the Techno-optimists Are Wrong'. *Foreign Affairs*. 94(4). July/August.

Anna Breman is Group Chief Economist at Swedbank since 2015, while also serving as independent expert in the government analysis group “The future of Work,” holding a Ph.D in Economics from Stockholm School of Economics. She previously worked as deputy director at the Swedish Ministry of Finance, and as a researcher in the US.

Diginomics and the productivity puzzle

Introduction

Technological change is the key driver of productivity growth¹. And productivity growth, in combination with total hours worked, determines GDP growth. As western countries face aging populations and weak or falling hours worked, productivity will determine GDP per capita and long-run economic growth. As Paul Krugman famously stated, “Productivity is not everything, but in the long run it is almost everything”.

It’s puzzling why we are not seeing higher productivity growth given recent digital innovations. Data shows a clear downward shift in productivity growth in developed countries since the mid-2000s. This slowdown in productivity growth pre-dates the financial crisis and therefore cannot be solely attributed to lower investment in the post-crisis period². There is a vibrant international discussion on ‘secular

1. Philippe Aghion and Peter Howitt, *Capital, innovation, and growth accounting*, Oxford Review of Economic Policy, Vol. 23, No. 1, 2007, pp.79–93

2. John G. Fernald, *Productivity and Potential Output Before, During, and After the Great Recession*, in Jonathan A. Parker and Michael Woodford, eds., NBER Macroeconomics Annual, Vol. 29 2014.

stagnation', low economic growth with stagnant productivity growth³. Unless productivity growth picks up, the theory posits that global growth will stay stagnant for the foreseeable future. With populist movements growing on the back of years of economic stagnation, weak productivity growth is a subject that should be at the forefront of the economic and political agenda.

This chapter presents data on productivity growth and investigates various explanations for the productivity puzzle. Based on existing literature, the explanations for the productivity puzzle are categorised into three main theories. First, the lack of productivity growth in the past has caused some researchers to say digitalisation is hype. They claim digital innovation is not comparable to other significant technological breakthroughs, which is why it has not led to faster productivity growth⁴. Second, another view claims that digitalisation causes productivity to be poorly measured, and more specifically, underestimated⁵. This could mean that we underestimate the recent increase in welfare due to new technologies. Third, economists and economic historians point out that, historically, there has always been a significant time lag between new innovations and their effect on productivity⁶. Hence, digital innovation presents a great potential to spur higher productivity in the future. We are simply not there yet.

In addition to these three main theories, productivity will be affected by the composition of the workforce and the relative size of

3. See e.g., Lawrence Summers, U.S. *Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound*, Business economics, Vol. 49, No. 2, 2014, pp. 65-73; Gauti B. Eggertson et al., *Secular Stagnation in the Open Economy*, American Economic Review, vol 106(5), 2016, pp. 503-507; Joel Mokyr, 'Secular Stagnation? Not in Your Life' in Coen Teulings and Richard Baldwin, eds., *Secular Stagnation: Facts, Causes, and Cures* VoxEU.org eBook, London: CEPR Press, 2014, pp. 83-89.

4. Robert J. Gordon, *The Rise and Fall of American Growth*, Princeton and Oxford, Princeton University Press, 2016; Paul Krugman, *The Big Meh*, The New York Times, 25.5.2015.

5. Philippe Aghion and Xavier Jaravel, *Knowledge Spillovers, Innovation and Growth*, the Economic Journal, Volume 125, Issue 583, pages 533-573, 2015; Erik Brynjolfsson and Andrew McAfee, *The Second Machine Age: Work Progress, and Prosperity in a Time of Brilliant Technologies*, New York: W. W. Norton & Company, 2014; Joel Mokyr, *Big Ideas - Riding the Technology Dragon*, Milken Institute Review, Milken Institute, Second Quarter 2014, pp. 8-94.

6. Joel Mokyr, *The Next Age of Invention: Technology's future is brighter than pessimists allow*, in Brian C. Anderson ed., City Journal, Manhattan Institute, Winter 2014, pp. 14-20.; Martin Fleming, *Digitization changes everything: improving economic measurement in an era of radical innovation and transformation*, In Monthly Labor Review, October 2015; Eric J. Bartelsman, *ICT, Reallocation and Productivity*, European Economy. Economic Papers 486, April 2013. Brussels. PDF, 44pp.

different sectors in the economy⁷. Digitalisation, automation and robotisation are changing the structure of the labour market. As routine tasks are being replaced, many workers find themselves working in the service sector where productivity, and hence wages, are traditionally low. This compositional effect may be important for overall productivity growth.

This paper focuses primarily on the direct link between technology and productivity as an explanation for the productivity puzzle. I argue there is merit to both the measurement problem and to the hype theory. Ultimately, I find the lag theory the most convincing of the three main explanations for the productivity puzzle. As a consequence, it is reasonable to expect, even though the effects of digital innovation on productivity growth so far have been lacklustre, that there is an upside to productivity growth in the years to come.

Rapid innovation, yet weak productivity growth


Productivity is measured as output per hour worked. There are many ways such efficiency gains can play out in different sectors. For instance, in manufacturing, digitally connected goods can free labour resources by automatically sending information through the production chain. Within the transport sector, autonomous vehicles can significantly increase output per hour worked. And, in the retail sector, e-commerce reduces hours worked in traditional retail sales. More specifically, digitalisation can transform innovation through: 1) improved measurement of business activities, 2) faster and cheaper business experimentation, 3) easier collaboration and sharing of ideas, and; 4) better ability to replicate processes and product innovations⁸. In combination, these aspects create a new kind of information and communication technology (ICT)-fueled R&D that could, with the right support, span the

7. See e.g., Andy Feng and Georg Graetz, *Rise of the Machines: The Effects of Labor-Saving Innovations on Jobs and Wages*, CEP Discussion Papers dp1330, Centre for Economic Performance, LSE, 2015; Anna Breman, *Diginomics och arbetet i framtiden*, Underlagsrapport till analysgruppen Arbetet i framtiden, 2015.

8. Erik Brynjolfsson, *ICT, Innovation and the E-Economy*, in EIB Papers Volume 16, No. 2/2011, 2011, pp. 60-76.

entire economy. The positive potential for digital technologies to increase productivity and boost economic growth is significant.

Figure 1: GDP growth, productivity and total hours worked, Sweden 1980-2015. Source: OECD.




Data, however, paints a grim picture of actual productivity growth. In the past decade, productivity growth has slowed significantly in developed countries. This slowdown occurred ahead of the financial crisis and cannot solely be explained by lower investments due to the Great Recession⁹. The slowdown is in stark contrast to the development during the productivity boom fueled by information and communications technologies during the 1990s and early 2000s. For example, Swedish productivity growth fell from 2.8 % per year during 1995 to 2005 to a very weak 0.45% per year in 2006-2015. The equivalent difference in productivity growth rates for the US is equally staggering: productivity growth was at 2.37% per year in 1995-2005 and 0.98 % in 2006-2015¹⁰.

9. John G. Fernald, *Productivity and Potential Output Before, During, and After the Great Recession*, in Jonathan A. Parker and Michael Woodford, eds., NBER Macroeconomics Annual, Vol. 29 2014.; John G. Fernald and Charles J. Jones, *The Future of U.S. Economic Growth*, Federal Reserve of San Francisco Working Paper 2014-02, 2014; Byrne et al, *Does the United States have a Productivity Slowdown or a Measurement Problem?* Federal Reserve Bank of San Francisco Working Paper 2016-03, 2016.

10. Calculations on OECD data.

As a result of this weak productivity growth, overall economic growth has been stagnant. This has led to a broad discussion on secular stagnation¹¹ and the end of economic growth¹².

Figure 2: The difference between trend productivity and actual productivity, Sweden 2006-2015. Source: OECD.



The productivity puzzle – three possible explanations

Weak productivity growth has also called into question the optimistic view on digitalisation and its potential to contribute to economic growth. If digitalisation is so transformative, how come productivity growth is so poor? Let us divide the explanations into three categories; (1) the hype theory; (2) the mismeasurement theory, and; (3) the lag theory.

11. Gaulti B. Eggertson et al, *A Contagious Malady? Open Economy Dimensions of Secular Stagnation*, NBER Working Paper No. 22299, 2016; Gaulti B. Eggertson et al., *Secular Stagnation in the Open Economy*, American Economic Review, vol. 106(5), 2016, pp. 503-507, Lawrence Summers, U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound, Business economics, Vol. 49, No. 2, 2014, pp. 65-73.

12. Robert J. Gordon, *Is US Economic Growth Over? Faltering Innovation Confronts the Six Headwinds*. NBER Working paper 18315, 2012; Robert J. Gordon, *The Rise and Fall of American Growth*, Princeton and Oxford, Princeton University Press, 2016.

(1) The hype theory

The most renowned among the critics of the ‘digital revolution’ has been Professor Robert Gordon, who argues that digitalisation is not comparable to the main technological breakthroughs of the past. In his book *The Rise and Fall of American Growth*, Gordon argues that the 10 decades between 1870 and 1970 represented the most rapid period of growth in labour productivity experienced in American history. This period brought a fundamental change in many dimensions of human life by producing many transformative innovations, the likes of which could never be invented again. The inventions after 1970 were indeed revolutionary within their domain – entertainment, communications and information technology. However, they “did not have the same effects on living standards as had electricity, the combustion engine, running water.”¹³

Gordon expects a permanent downward shift in productivity growth in the future. He does not consider digital innovations such as 3D printing, driverless cars, robotics and AI to have the potential to increase productivity growth to the same extent as the ICT-induced productivity boost during the 1990s and early 2000s. Gordon heavily emphasises data and historical evidence from the past decade to make this argument.

As a comment to the debate between the techno-optimist and the sceptics Paul Krugman wrote an op-ed on digitalisation in the *New York Times* saying “We ought to scale back the hype.”¹⁴ As many others, Krugman did not say that digitalisation cannot spur productivity growth, but cautioned that the potential may not be as transformative as many would like to believe.

(2) The mismeasurement theory

In response to Gordon and Krugman, several researchers have point-

13. Robert J. Gordon, *The Rise and Fall of American Growth*, Princeton and Oxford, Princeton University Press, 2016.

14. Paul Krugman, *The Big Meh*, *The New York Times*, 25.5.2015, <http://www.nytimes.com/2015/05/25/opinion/paul-krugman-the-big-meh.html> (accessed 9 September 2016).

ed out that measuring productivity growth poses challenges and that these challenges have increased as digital inventions have taken off.

Measuring output and productivity is complex and ambiguous and has a long history of controversies, dating back to the introduction of national accounts¹⁵. Lately, these discussions have been revived by the emergence of digital innovations and the puzzling fact that seemingly revolutionary technologies do not show up in productivity statistics. The proponents of the mismeasurement theory have three key arguments; firstly, free goods and services do not show up in the measurement of GDP. Hence digital services that are free for consumers, such as Facebook, LinkedIn and Google Maps, are not included in measurements of GDP. Second, digital products rely more heavily on intangible capital, which is harder to measure. This includes immaterial inputs such as patents and copyright, as well as human capital. Third, quality-adjusted prices are difficult to measure for digital services. For example, streaming digital content such as music and movies vastly increases quantities available to consumers at a fixed cost¹⁶.

Brynjolfsson and Oh¹⁷, and Brynjolfsson and McAfee¹⁸ therefore argue that the increased consumer surplus due to an increase in the supply of free digital services should be taken into account in official statistics. Others claim that the benefits of technological progress have in the past been overstated as negative externalities, such as climate change and pollution, and were not taken into account¹⁹.

15. Diane Coyle, *GDP: A Brief But Affectionate History*, Princeton, NJ: Princeton University Press, 2014; Joel Mokyr, *Big Ideas – Riding the Technology Dragon*, Milken Institute Review, Milken Institute, Second Quarter 2014, pp. 8-94.

16. Erik Brynjolfsson and Andrew McAfee, *The Second Machine Age: Work Progress, and Prosperity in a Time of Brilliant Technologies*, New York: W. W. Norton & Company, 2014, Martin Fleming, *Digitization changes everything: improving economic measurement in an era of radical innovation and transformation*, In Monthly Labor Review, October 2015, <http://www.bls.gov/opub/mlr/2015/article/digitization-changes-everything-1.htm> (accessed 9 September 2016), Joel Mokyr, 'Secular Stagnation? Not in Your Life' in Coen Teulings and Richard Baldwin, eds., *Secular Stagnation: Facts, Causes, and Cures* VoxEU.org eBook, London: CEPR Press, 2014, pp. 83-89.

17. Erik Brynjolfsson and Joo Hee Oh, *The Attention Economy: Measuring the Value of Free Digital Services on the Internet*, International Conference on Information Systems, ICIS 2012, 2012, 4:3243-3261

18. Erik Brynjolfsson and Andrew McAfee, *The Second Machine Age: Work Progress, and Prosperity in a Time of Brilliant Technologies*, New York: W. W. Norton & Company, 2014.

19. Joel Mokyr, *Big Ideas – Riding the Technology Dragon*, Milken Institute Review, Milken Institute, Second Quarter 2014, pp. 8-94.

In an attempt to correct for mismeasurement problems in the face of new technology, Byrne et al (2016) conduct a time series analysis using three alternative price indices for the IT sector and compare the result to national income and products accounts (NIPAs) estimates on IT prices to estimate the mismeasurement of official accounts²⁰. The authors find that mismeasurement of IT products was significant already before the slowdown of productivity growth and that there is no evidence that it has gotten worse. In addition, the authors calculate the increase in consumer surplus in the non-market sector. They find that these gains, although not insignificant, are far from large enough to compensate from the slower growth in the market sector.

The literature on mismeasurement does highlight that there are indeed real difficulties in the measuring digital services. The topic warrants more research, and further study is in progress²¹. The conclusion, however, cannot be that mismeasurement explains the full reduction in productivity gains. Even if higher consumer surplus is included, productivity and economic growth has slowed down significantly since mid-2000s.

(3) The lag theory

The question remains whether digital innovations in combination with higher investment can spur productivity growth, thus bringing western economies back to growth levels seen before the financial crisis. Is there simply a time lag between innovations and productivity gains, which means that the best is yet to come? ‘The Next Age of Innovation: Technology’s future is brighter than pessimists allow’ is the title of a paper by Joel Mokyr²² arguing that the best is indeed yet to come. Pessimistic forecasts about the end of important innovations have turned out to be false before. Technological progress enhances ac-

20. Byrne et al, *Does the United States have a Productivity Slowdown or a Measurement Problem?* Federal Reserve Bank of San Francisco Working Paper 2016-03, 2016.

21. Aghion et al, *Missing Growth From Creative Destruction*, Working paper September 2016.

22. Joel Mokyr, *The Next Age of Invention: Technology’s future is brighter than pessimists allow*, in Brian C. Anderson ed., *City Journal*, Manhattan Institute, Winter 2014, pp. 14-20.

cess to already existing knowledge and information which helps scientific progress. A number of authors point to the potential of digital innovation to transform larger sectors of the economy such as health care, transport and energy²³.

What do we know about time lags and technological change? Historically, eras of rapid innovations have been associated with relatively low productivity growth. It takes time for technologies to diffuse into common use. The benefits of productivity-enhancing but disruptive technologies have been reaped only slowly, with an initial lag, over the course of decades. For example, the 1930s did see rapid inventions, while the highest levels of productivity growth were seen only in the 1940s and the 1950s as these new technologies adopted into use in production²⁴. In 1987, the Nobel-winning economist Paul Samuelsson famously stated, “you can see the computer age everywhere except in the productivity statistics”. It took until the mid-1990s until the computer age and IT started to produce high productivity growth.

In addition, some authors suggest that we can see eras of new innovation being associated with temporary productivity slumps²⁵. This argument goes beyond the time lag explanation by hypothesising that periods of rapid innovation can actually see lower productivity gains due to adaptation costs of new technology. Restructuring of workplaces can be slow and costly as old and new technologies tend to be used side by side, causing a downward shift in productivity for some time. Eventually, workplaces and production methods are fully adapted to the new technology and productivity growth will again increase.

23. Fleming et al., *Economic Vitality 2.0: Prosperity and Public Engagement in a Data-Driven World*, IBM Public Sector Thought Leadership Whitepaper, 2015; David Rotman, *Dear Silicon Valley: Forget Flying Cars, Give Us Economic Growth*, MIT Technology Review, 21.6.2016. Available at: <https://www.technologyreview.com/s/601682/dear-silicon-valley-forget-flying-cars-give-us-economic-growth/> (accessed 9 September 2016); Bloom et al, *Management as Technology?*, CEPR Discussion Paper No. DP11312, 2016; Ollivaud et al, *Links between weak investment and the slowdown in productivity and potential output growth across the OECD*, Paris, OECD Economics Department Working Papers, No. 1304, OECD Publishing, 2016; Erik Brynjolfsson, *ICT, Innovation and the E-Economy*, in EIB Papers Volume 16, No. 2/2011, 2011, pp. 60-76; Eric J. Bartelsman, *ICT, Reallocation and Productivity*, European Economy. Economic Papers 486, April 2013. Brussels. PDF, 44pp.

24. Robert J. Gordon, *The Rise and Fall of American Growth*, Princeton and Oxford, Princeton University Press, 2016.

25. For an overview, see Krusell, Per, 2000: “Ny Teknik och Produktivitet – Vad Vet Vi?”. *Ekonomisk Debatt* 2000: 28:6.

Will we enter the second wave of the IT era with higher productivity gains to come in the near future? Digitalisation is transforming society, the labour market and the economy in profound ways²⁶. Nevertheless, its productivity-enhancing possibilities are not yet very impressive. Digital technologies have had the greatest impact on a limited set of sectors, such as music, entertainment and journalism. Free digital services, such as Facebook, LinkedIn and Google Maps may indeed increase the consumer surplus and transform communication, social networks and the diffusion of knowledge and ideas, but they do not vastly increase output per hour worked above and beyond what the internet has done already.

So far, mobile internet only marginally increases productivity compared to having the internet through a fixed land-line. For example, banking services, such as paying bills, can now easily be done on a mobile phone. This may enhance consumer surplus by making it more convenient, but it increases efficiency only marginally compared to paying bills on a computer. From the business side, however, a consumer performing banking services that previously required contact with staff at a bank does have the potential to vastly increase productivity in financial services, regardless of whether this is done on a mobile or stationary device. The large-scale effect of that transformation is still ahead of us. The financial sector could very well be in the adaptation phase described by Krusell (2000), where old and new technologies are being used side-by-side and the full productivity gains will be made only when the transformation is complete.

I find the lag theory the most convincing of the three possible explanations for the productivity puzzle. Current weak productivity growth is likely to stem from a combination of factors, where adaptation costs of new technology is just one. Weak investment following the financial crisis, further slowing the adaptation is another. The third is a compositional effect in the labour market, where people displaced by new la-

26. Anna Breman and Anna Felländer, *Diginomics – Nya Ekonomiska Drivkrafter*, Ekonomisk Debatt, nr. 6 2014, årgång 42; Anna Breman, *Diginomics och arbetet i framtiden*, Underlagsrapport till analysgruppen Arbetet i framtiden, 2015.

bour-saving innovations are entering the less-productive, but more labour-intensive, service sector.

Measurement problems do indeed exist, but there is no compelling evidence that they can explain the entire slowdown in productivity growth. There is also merit to the hype theory. The inventions of the industrial revolution did transform production capacity in an unparalleled way, something that the digital technologies may not be able to fully match. There is, nonetheless, much catching up to do simply to reach the existing technological frontier. In addition, there are inefficiencies in large sectors in advanced economies, such as transport, energy and health care. As a consequence, there is scope for new inventions. It is reasonable to expect, even though the effects of digital innovation on productivity growth so far have been lacklustre, that there is an upside to productivity growth in the years to come.

Conclusions

The world economy is suffering from weak growth in the aftermath of the financial crisis. In many western countries, households have seen stagnant real wages and weak labour markets. In this environment of prolonged weak growth, digital innovations that spur productivity growth can contribute to higher growth, increased real wages and thus brighter prospects for long-run wellbeing.

Research on the productivity paradox may provide some insights into the potential to conduct pro-growth policies. If digitalisation is merely hype, the room for policy to improve productivity growth is limited. If we cannot measure productivity properly, there is no paradox, simply a mismeasurement problem and no need for pro-growth policies. If there is a time lag between inventions and productivity growth, it may be possible to improve productivity growth through economic policy. In particular, it may be possible reduce the time lag between innovation and the adaption of new technologies through, for example, regulatory change and public investment.

The slowdown in productivity growth since the mid-2000s has been exacerbated by the Great Recession which has reduced capital availa-

bility. The OECD finds that relatively low capital investments in both tangible and intangible goods have severely dampened the diffusion and adaptation of digital inventions²⁷. As demand remains weak and the outlook for economic growth has stagnated, firms have been reluctant to invest despite low interest rates. Higher demand through, for example, fiscal stimulus, could potentially increase aggregate demand and increase business confidence and investment²⁸. The large productivity gains will come from digital innovation contributing to higher productivity within a broader set of sectors in the economy. Sectors, such as health care, energy and transportation, may be more capital-intensive compared with the sectors that have already been transformed, such as media and entertainment. A combination of regulatory changes and public investment may therefore help accelerate the transformation of these sectors.

Furthermore, a large body of research finds that education is key to fostering R&D into new inventions, and to foster adaptation of new technologies²⁹. The growth in educational attainment that helped drive productivity growth through the 20th century has levelled off³⁰. Policies to strengthen educational attainment can speed the recovery of productivity growth. Socioeconomic factors are still highly correlated to educational attainment in many countries, which is likely to cause an inefficient utilisation of human capital. In addition, rapid technological change warrants reforms that allow employees to more actively engage in educational activities throughout their working life³¹. Life-long learning is likely to become increasingly important in order to keep up with

27. Ollivaud et al, *Links between weak investment and the slowdown in productivity and potential output growth across the OECD*, Paris , OECD Economics Department Working Papers, No. 1304, OECD Publishing, 2016.

28. Lawrence Summers, *U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound*, Business economics, Vol. 49, No. 2, 2014, pp. 65-73; Gaulti B. Eggertson et al., *Secular Stagnation in the Open Economy*, American Economic Review, vol 106(5), 2016, pp. 503-507.

29. Claudia Goldin and Lawrence F. Katz, *The Race between Education and Technology*, Cambridge, Harvard University Press 2008; Robert J. Gordon, *The Rise and Fall of American Growth*, Princeton and Oxford, Princeton University Press, 2016.

30. Claudia Goldin and Lawrence F. Katz, *The Race between Education and Technology*, Cambridge, Harvard University Press 2008; Robert J. Gordon, *The Rise and Fall of American Growth*, Princeton and Oxford, Princeton University Press, 2016.

31. Anna Breman, *Diginomics och arbetet i framtiden*, Underlagsrapport till analysgruppen Arbetet i framtiden, 2015.

technological change. Universities could be incentivised to provide a wider range of courses that allow retraining and life-long learning.

I believe new innovations and faster adaptation of existing technologies have the potential to enable productivity increases similar to those during the 1990s IT boom. The lag theory does provide plausible explanations for why productivity growth has tended to come and go throughout the last century. Bartelsman estimates that ICT could enable growth to rise back to an annual rate of 2.5%³². However, I am very skeptical as to whether we will actually see higher growth in productivity trends in the coming years. There is very little concerted effort to making regulatory changes to spur innovation and investments in digital technologies. In addition, political and economic constraints hold back both public and private investments. In the short term, my expectation is for continued weak productivity growth. Trend growth of around 1% per year, or less, is a more likely scenario, as compared to trend growth returning to 2.5% per year in advanced economies. Long term, however, I am cautiously optimistic.

32. Eric J. Bartelsman, *ICT, Reallocation and Productivity*, European Economy. Economic Papers 486, April 2013. Brussels. PDF, 44pp.

References

Aghion, Philippe and Marinescu, Ioana, 2011: 'Cyclical Budgetary Policy and Economic Growth: What Do We Learn from OECD Panel Data?' In Phelps, Edmund S. and Sinn, Hans-Werner (ed), 2011: 'Perspectives on the Performance of Continental Economics.' CESifo Seminar Series, MIT Press.

Aghion, Philippe and Howitt, Peter, 2007: 'Capital, innovation, and growth accounting' In Oxford Review of Economic Policy, Volume 23, Number 1, 2007, pp.79–93.

Aghion, Philippe; Howitt, Peter and Prantl Susanne, 2013: 'Patent Rights, Product Market Reforms and Innovation' NBER Working Paper No. 18854.

Aghion Philippe; Akcigit, Ufuk and Howitt, Peter, 2013: 'What do we learn from Schumpeterian Growth Theory?' PIER Working Paper No. 13-026.

Aghion, Philippe and Jaravel, Xavier, 2015: 'Knowledge Spillovers, Innovation and Growth' In the Economic Journal, Volume 125, Issue 583, pages 533–573, March 2015. Available at: <http://onlinelibrary.wiley.com/doi/10.1111/eoj.12199/epdf>.

Aghion et al, 'Missing Growth From Creative Destruction', Working paper, September 2016.

Bank for International Settlements, 2016: '86th Annual Report.' Bank for International Settlements. Available at: www.bis.org/publ/arpdf/ar2016e.htm.

Bartelsman, Eric J., 2013: 'ICT, Reallocation and Productivity' In European Economy. Economic Papers 486. April 2013. Brussels. PDF. 44pp. Available at: http://ec.europa.eu/economy_finance/publications/economic_paper/2013/ecp486_en.htm.

Bloom, Nicholas; Sadun, Raffaella and Van Reenen, John, 2016: 'Management as Technology?' CEPR Discussion Paper No. DP11312. Available at: http://cep.lse.ac.uk/_new/publications/abstract.asp?index=5068.

Breman, Anna and Felländer, Anna, 2014: 'Diginomics – Nya Ekonomiska Drivkrafter' *Ekonomisk Debatt*, nr. 6 2014, årgång 42. Available at: <http://nationalekonomi.se/sites/default/files/NEFfiler/42-6-abaf.pdf>.

Breman, Anna, 2015: 'Diginomics och arbetet i framtiden' Underlagsrapport till analysgruppen Arbetet i framtiden. Available at: <http://www.regeringen.se/contentassets/17116303997443d69a6692270d963b01/underlagsrapport-diginomics-och-arbetet-i-framtiden.pdf>.

Brynjolfsson, Erik, 2011: 'ICT, Innovation and the E-Economy' In EIB Papers Volume 16. n°2/2011. Available at: <http://www.eib.org/infocentre/publications/all/eibpapers-2011-v16-no2.htm>.

Brynjolfsson, Erik and McAfee, Andrew, 2014: 'Beyond GDP: How Our Current Metrics Mismeasure the Digital Economy.' In Brynjolfsson, Erik and McAfee, Andrew, 2014: *The Second Machine Age: Work Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W. W. Norton & Company, 2014. Available at: <http://www.markle.org/sites/default/files/Beyond%20GDP.pdf>.

Brynjolfsson, Erik and Oh, Joo Hee, 2012: 'The Attention Economy: Measuring the Value of Free Digital Services on the Internet.' International Conference on Information Systems, ICIS 2012, 2012, 4:3243-3261. Available at: <http://aisel.aisnet.org/cgi/viewcontent.cgi?article=1045&context=icis2012>.

Byrne, David; Fernald, John and Reinsdorf, Marshall, 2016: 'Does the United States have a Productivity Slowdown or a Measurement Problem?' Federal Reserve Bank of San Francisco Working Paper 2016-03. Available at: <http://www.frbsf.org/economic-research/files/wp2016-03.pdf>.

Coyle, Diane, 2014: 'GDP: A Brief But Affectionate History.' Princeton, NJ: Princeton University Press.

Draghi, Mario, 2016: 'On the Importance of Policy Alignment to Fulfill Our Economic Potential.' European Central Bank, 2016. Available at: <https://www.ecb.europa.eu/press/key/date/2016/html/sp160609.en.html>.

Eggertson, Gauti B.; Mehrotra, Neil R.; Singh, Sanjay R. and Summers, Lawrence H., 2016: 'A Contagious Malady? Open Economy Dimensions of Secular Stagnation.' NBER Working Paper No. 22299. Available at: <https://sites.google.com/site/neilmehrotra/>.

Eggertson, Gauti B.; Mehrotra, Neil R. and Summers, Lawrence H., 2016: 'Secular Stagnation in the Open Economy.' *American Economic Review*, vol 106(5), pages 503-507. Available at: <https://sites.google.com/site/neilmehrotra/>.

Feng, Andy & Graetz, Georg, 2015: 'Rise of the Machines: The Effects of Labor-Saving Innovations on Jobs and Wages,' CEP Discussion Papers dp1330, Centre for Economic Performance, LSE. Available at: <https://ideas.repec.org/p/cep/cepdps/dp1330.html>.

Fernald, John G. and Jones, Charles J., 2014: 'The Future of U.S. Economic Growth.' Federal Reserve of San Francisco Working Paper 2014-02. Available at: www.frbsf.org/economic-research/files/wp2014-02.pdf.

Fernald, John G., 2014: 'Productivity and Potential Output Before, During, and After the Great Recession.' In *NBER Macroeconomics Annual*, Vol. 29, ed. by Parker and Woodford, 2014. Available at: <http://www.frbsf.org/economic-research/economists/john-fernalld/>.

Fleming, Martin, 2015: 'Digitization changes everything: improving economic measurement in an era of radical innovation and transformation' In *Monthly Labor Review*, October 2015. Available at: <http://www.bls.gov/opub/mlr/2015/article/digitization-changes-everything-1.htm> (accessed 9 September 2016).

Fleming, Martin; Dencik, Jacob and Forcke Anne-Rivers, 2015: 'Economic Vitality 2.0: Prosperity and Public Engagement in a Data-Driven World' IBM Public Sector Thought Leadership Whitepaper. Available at: <http://www.lisboncouncil.net/news-a-events/644.html>.

Goldin, Claudia and Katz, Lawrence F., 2008: 'The Race between Education and Technology', Harvard University Press, Cambridge.

Gordon, Robert J., 2012: 'Is US Economic Growth Over? Faltering Innovation Confronts the Six Headwinds.' NBER Working paper 18315. Available at: <https://ideas.repec.org/p/nbr/nberwo/18315.html>.

Gordon, Robert J., 2016: 'The Rise and Fall of American Growth.' Princeton University Press, Princeton and Oxford.

Krugman, Paul, 2015: 'The Big Meh.' The New York Times, 25.5.2015. Available at: <http://www.nytimes.com/2015/05/25/opinion/paul-krugman-the-big-meh.html> (Accessed 9 September 2016).

Krusell, Per, 2000: 'Ny Teknik och Produktivitet – Vad Vet Vi?'. Ekonomisk Debatt 2000: 28:6.

Mokyr, Joel: 'Big Ideas – Riding the Technology Dragon' in Milken Institute Review, Milken Institute, Second Quarter 2014, pp. 8-94. Available at: <http://sites.northwestern.edu/jmokyr/research/>.

Mokyr, Joel, 2014: 'The Real Future of Capitalism/Capitalism Reinvents Itself' in Current History, History of Capitalism, Vol. 112, Issue 757, 2013, pp. 291-297. Available at: <http://sites.northwestern.edu/jmokyr/research/>.

Mokyr, Joel, 2014: 'Secular Stagnation? Not in Your Life.' In Coen Teulings and Richard Baldwin, eds., 2014: 'Secular Stagnation: Facts, Causes, and Cures.' VoxEU.org eBook, London: CEPR Press, 2014, pp. 83-89. Available at: <http://sites.northwestern.edu/jmokyr/research/>.

Mokyr, Joel, 2014: 'The Next Age of Invention: Technology's future is brighter than pessimists allow.' in Brian C. Anderson ed., City Journal, Manhattan Institute, Winter 2014, pp. 14-20. Available at: <http://sites.northwestern.edu/jmokyr/research/>.

OECD, 2016: 'Employed population, Aged 15 and over, All persons', Short-Term Labour Market Statistics (database) [Online]. Available at: <http://stats.oecd.org/Index.aspx?DataSetCode=QNA> [Accessed on 28 July 2016].

OECD, 2016: 'GDP per capita and productivity growth', OECD Productivity Statistics (database) [Online]. Available at: DOI: <http://dx.doi.org/10.1787/data-00685-e> [Accessed on 28 July 2016].

OECD, 2016: 'Gross domestic product _ expenditure approach', Quarterly National Accounts (database) [Online]. Available at: <http://stats.oecd.org/Index.aspx?DataSetCode=QNA> [Accessed on 28 July 2016].

OECD, 2016: 'Average annual hours actually worked per worker', LFS - Average Annual Hours Worked (database) [Online]. Available at: <https://stats.oecd.org/Index.aspx?DataSetCode=ANHRS> [Accessed on 28 July 2016].

Ollivaud, Patrice; Guillemette, Yvan and Turner, David, 2016: 'Links between weak investment and the slowdown in productivity and potential output growth across the OECD', OECD Economics Department Working Papers, No. 1304, OECD Publishing, Paris.

Rotman, David, 2016: 'Dear Silicon Valley: Forget Flying Cars, Give Us Economic Growth.' MIT Technology Review, 21.6.2016. Available at: <https://www.technologyreview.com/s/601682/dear-silicon-valley-forget-flying-cars-give-us-economic-growth/>.

Summers, Lawrence H., 2014: 'U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound.' In Business economics, Vol. 49, No. 2. Available at: <http://larrysummers.com/wp-content/uploads/2014/06/NABE-speech-Lawrence-H.-Summers1.pdf>.

Fredrik Löfgren works at Linköping University as a roboticist, he also gives lectures about robotics and develops robots in his own company. He has been constructing robots for over 15 years.

How may robots affect the labour market in the near future?¹

Introduction

This chapter discusses how different applications for robots will affect the labour market in the near future. *Near future* refers to the next 10–50 years. It is likely that several occupations will disappear, but new ones will also emerge. However, we claim that the net result will be negative, which means that we will have higher unemployment. These effects will not happen overnight, and not all occupations will be affected. But, this will happen for a sufficient amount of the population for it to become a problem for society.

The observations made in this chapter are not from the point of view of a social scientist, but that of a roboticist. The observations are taken together with readings of scientific literature on automation. I do not claim to have answers to the economic and social scientific problems thrown up, but to raise a set of critical questions for the reader.

All the examples in this chapter are real technologies that exist, not just in science-fiction or future technology. However, most of the examples are still in their research stage and are either not available for the general public, or still very expensive.

1. Thanks to Jennifer Krieger, Karl Wennberg and the reviewers for comments and suggestions on the text.

No one can predict the *future in detail*, but this chapter tries to provide a scenario of the future of different kinds of occupations through the perspective of the field of robotics. I have been developing robots for 15 years and will use some examples that I have constructed, but also examples from other roboticists. The chapter does not discuss the risks of automation for all occupations, but instead focuses on blue-collar workers, such as machine operators, the transportation sector with the advent of driverless cars, white-collar workers in offices, skilled professions in the legal and medical spheres, and creative workers.

This chapter raises the following questions:

- Will we have a job to go to in the near future?
- How can we earn our living when more and more jobs will be done by robots?
- How will our economic system handle the increased unemployment?

The objective of this chapter is to raise awareness of potential risks with robotisation, and to spur thinking about how society may change in the near future.

What is a robot?

The *Oxford English Dictionary* define *Robot* as:

“A machine capable of carrying out a complex series of actions automatically, especially one programmable by a computer.”

Robots can be categorised in two ways: physical robots, such as industrial robots, and immaterial robots (software or algorithms), such as stock exchange robots. Each copy of a physical robot must be manufactured, while software robots are just source code that is easy to copy. Software robots are cheaper to reproduce and duplicate than physical robots. However, this doesn't reflect the development costs

of a robot. Both physical and immaterial robots need maintenance. Also, physical robots are not just a mechanical structure, but a combination of hardware and software working together. Another classification for robots is based on their level of autonomy, i.e. how much the robot can do by itself. An industrial robot has a low level of autonomy, since it is just following predefined movement instructions and is not aware of its environment. Conversely, a driverless car is much more aware of its environment and therefore has high autonomy. A robotic lawn mower is somewhere in between, because it has a very basic awareness of its surroundings.

The word *robot* was first used by Karel Čapek 1921 in a play called *Rossumovi Univerzální Roboti*, meaning *Rossum's Universal Robots*. Karel described artificial life that was constructed by humans as workers for humanity. The word derives from the Old Church Slavic *rabota* meaning *servitude*.

Background

When *Homo sapiens* first appeared in the African savannah, almost all our time was spent as hunter-gatherers: hunting, collecting roots and fruits, and fishing. But humans are lazy by nature². We do not want to spend all our time gathering food. Instead we created solutions to help us, for example archery and domesticated animals. This has continued through history, evolving into ever-more complicated solutions.

Humans have created machines since ancient times. Ancient Egyptians used complicated irrigation systems and mechanical solutions to build the pyramids. The Greek mathematician Heron wrote and experimented with hydraulics, mechanics, fire engines and even programmable carts.

2. Mar-Nicolle, 'Wired for Laziness'; Hreha, 'Is Human Laziness Human Nature?'

The industrial revolution 2.0?

At the end of 1700s, 90% of the population in the US was working as farmers³. This fell to below 1.7% in 2011.⁴ Over the same period the US population has increased from 3.9 million to over 300 million.⁵ The industrial revolution enabled humans to replace our efforts with a mechanical counterpart. Not because our muscles were weaker or no longer functioning, but because mechanical muscle were far more powerful and durable. Humans did not change, instead the mechanical solutions we invented surpassed us.

The industrial revolution changed the foundations of the labour market. Farming jobs disappeared when tractors and threshing machines were introduced. At the same time, new occupations were formed in the new industrial sector, such as welders and assembly-line workers.

Society changed and people changed with it. More people moved from the countryside into cities to work in the new industries. New cities emerged around these new industries. The efficiency of production was increased when companies and people became more specialised, though this often requires more educated workers. This led to public schooling and countries introduced compulsory education. With education, people became more aware of their rights, and argued for having rights. Labour unions were formed. Working conditions have improved over the last 100 years, with the standard of living today considerably higher than before the industrial revolution. We have better working conditions, higher salaries and thanks to new technology, better infrastructure and cities.

Now we are facing a new revolution, the digital revolution. This may be considered as a continuation of the industrial revolution, or as a new revolution that will affect society in a similar way. There are both similarities and differences to the industrial revolution. Occupations will disappear, just as some did during the 17th and 18th century. Society became more connected with roads and cars, and will continue to become

3. Diamandis and Kotler, *Abundance: The Future Is Better than You Think*.

4. Ibid.

5. 'Demographic History of the United States.'

How may robots affect the labour market in the near future?

even more connected. For example, the internet enables us to participate in meetings without being there physically. This enables people to move back to the countryside thanks to this new infrastructure⁶.

However, one big difference is that the digital revolution we are facing now is not just going to replace our muscles, but also part of our mind with a digital counterpart. We are already creating robots and programs with the ability to make decisions. For example, there are stock exchange robots that can act and make independent decisions.

This makes the impact on the society even bigger. What do humans have left when machines and robots both have muscles and decision-making capabilities? How can we compete on the labour market? Not all occupations will disappear, and neither will no new occupations be created. One recent example is the app developer occupation. 10 years ago they did not exist. In 2008 Apple launched the App Store and the mobile app industry exploded, generating \$41.1 billion of revenues in 2015⁷. So clearly new professions have emerged, and new technology has given us better ways to communicate and enhance our quality of life.

But the difference is that technology will continue to get better and better, while humans have barely changed since we migrated from the African savannah. Our biological evolution has no capacity to match the rate of development of computers and algorithms.

While the changes to our different occupations will not happen overnight, it has already started and will continue affecting society in the coming century. We must be prepared for this change and try to adapt ourselves and society to higher rates of unemployment.

Do new technologies lead to new jobs?

Horses have had different uses throughout history. They started out with labour-intensive tasks in mines, but were replaced with more-efficient and cheaper machinery. Humans found other tasks for horses, like carrying letters between cities. Horses live a better life now. Seen in this way, it is logical to think that new technology always leads to higher

6. Dominiczak, 'Countryside Population to Increase Dramatically by 2025.'

7. McGoogan, 'App Revenue Will Overtake the Music Industry This Year.'

standard of living.

There is no law of nature that states that new technology will create new and better professions. New technology leads to a higher standard of living and economic growth, but not necessarily to new professions. *Post hoc ergo propter hoc*, there is no causality between new and better technology and more jobs.

In fact when the new mechanical horses (i.e. cars) were invented at the end of 19th century, horses had a harder time finding a use. Horses were no longer profitable to the labour market and no longer paid for themselves. Cars were better suited for the new cities and infrastructure that emerged during the 20th century. The only occupation that horses are still today used is in the police and for forest conservation.

With our political leaders stating the need for investment in new technology, such as automation and robotisation, on the basis that it will create new job opportunities, in fact the opposite is true, the very outcome of automation is to reduce the demand for human jobs.

The argument that labour unions will ensure job opportunities stay in their home country, and prevent people from losing their jobs, has been proved by history not to endure in the long run, even if this may be true at the start. Labour unions and worker movements have always tried to prevent factory owners laying off personnel, but the economic incentives are so powerful and so strong that labour unions and worker movements eventually have been forced to succumb. For haulage contractor companies, personnel costs are the second-largest expense (the largest being diesel). Companies may try to reduce this expense to increase profit.

When James Hargreaves invented the Spinning Jenny in the 18th century, criticism was severe because many spinners were laid off and lost their only income. Textile workers formed groups, the Luddites, to sabotage the new machinery. The word *sabotage* originates from the French word for *clog* (*sabot*), when workers placed their shoes in the weaving looms. Also, the phrase *to put a spoke in someone's wheel* originates from the time of the industrial revolution. Labour unions have always tried to stop technology that threatens jobs. Yet after 5-10 years the Luddites disappeared (after military intervention), and some dec-

How may robots affect the labour market in the near future?

ades later workers eventually accepted the fact of new technology. In fact, the employee conditions have improved. Today, no one would like to change jobs with the dirty industrial environments of the late 18th century. But, it was hard for workers to give up their jobs to mechanical counterparts. The key point is that these changes were good for the society and future workers, because their conditions improved, but the changes were tough for the individual employees at the time.

In conclusion, new technology doesn't lead to new jobs; actually the goal of new technology is often the opposite, to automate a process. And even though individuals and groups such as labour unions sometimes try to prevent new technologies from replacing the human workforce, history shows they will eventually give up and accept the new conditions.

Blue-collar workers

The general public either associates the word *robot* with sci-fi robots, such as R2-D2, or with Industrial Robots that manufacture cars. Industrial robots are among the earliest robots that replaced the human workforce. The first commercial industrial robot was used in the 1960s and used in the automotive industry for dangerous tasks.

There are several advantages of using robots instead of humans in the manufacturing industries:

- A robot can execute heavy, monotonous and dangerous tasks, improving occupational health and safety (OHS) for humans.
- A robot is able to perform tasks repetitively and very precisely, leading to improved quality and decreased line rejection.
- A robot doesn't require any salary and is more efficient, leading to cheaper production.

All these factors lead to fewer employees in the manufacturing indus-

try. This has been the trend since the 1960s⁸. More products can be manufactured using fewer workers, and those that are working are undertaking supervision and maintenance instead of manufacturing. When my grandparents graduated from secondary school they could get a job at the ironworks or paper mill. Today you need a high level of education to work in the industry. It is no longer possible to just learn on the job as apprentice. We are getting rid of ‘easy’ tasks (skill level 1 and 2 on the ISCO-08 categorization⁹), and replacing them with robots, keeping the more advanced tasks (skill level 3 and 4¹⁰) for humans. This is a problem since not everyone has access to, or an interest in, further education.

In terms of development, new industrial robots are emerging, such as the collaborative robot Baxter from Rethink Robotics and YuMi from ABB. We can see an analogy in the history of computers. The industrial robots of the 1960s were much like the computers of the 1940s. The ENIAC computer didn’t have a user interface and no digital stored program memory, so cables had to be manually reconfigured. The first industrial robots took several months to reprogram for a new task. In the 1980s this changed, when computer experts started to develop home computers that were easy to use, programs could be switched using floppy disks and high-level programming languages were born, such as BASIC, which were able to reprogram computers (instead of reconnecting cables). Today, people can buy home robots, such as robotic vacuum cleaners and robotic lawnmowers, and the experts are building their own robots, just like they built computers in the 1980s. History repeats itself with an offset of 30 years.

In the 1980s, computer development exploded. We are likely to see the same happen to the robotics industry. When robots get easier to handle and don’t require expertise, the general public will start to use them in their everyday life. With demand increasing, the price will decrease. The speed and robustness of the robots will also improve, just like computers have increased in speed. 30 years after Apple introduced the Macintosh computer, everyone has an iPhone or other smartphone

8. Wyatt and Hecker, ‘Occupational Changes during the 20th Century.’

9. Hunter, *International Standard Classification of Occupations*.

10. *ibid.*

in their pocket that is several hundred times faster. Computers nowadays are also easier to interact with, using touch interfaces and voice command services like Siri.

The state-of-the-art robots are at the level where computers were in the 1980s. They are starting to get easier to use with lead-through techniques¹¹ and better programming tools (like ROS¹²), and are also starting to sneak into everyday use, i.e. lawnmowers. Collaborative robots are on the rise, with Rethink Robotics launching the Baxter in 2012¹³ and ABB releasing competitor YuMi in 2015¹⁴. Both these robots are intended to work side-by-side with humans. Baxter has eyes and a face to interact with his human co-worker and YuMi leans forward to indicate submission or subjection. These robots are often used for assembly of consumer electronics previously done by humans.

However, Baxter and YuMi are still very slow, much slower than their human counterparts. But they don't need coffee breaks, sleep or a salary, making them much cheaper operationally.

Today, robots are actually relatively cheap to acquire (or, if you like, to 'employ'¹⁵). Baxter costs about \$40,000¹⁶. The median salary for an assembler or welder in Sweden is approx. \$2,750/month¹⁷, with employment tax at \$850¹⁸, resulting in a cost of \$43,200 yearly. So, Baxter's return on investment is less than one year compared to a human doing the same work. While there are additional costs for programming Baxter, this is easy to do, and could be done using existing employees that knows the processes and production flow.

In 2014, I competed in an international cooking competition¹⁹, or rather, a robot I developed competed, since I can't cook at all. The competition was held in Madrid and the robot was taught to cook gazpacho.

11. Graf, Lead-through robot programming system.

12. 'Robot Operating System.'

13. 'Baxter (Robot).'

14. 'IRB 14000 YuMi.'

15. FinWire, 'EU-Utredning: Arbetsgivaravgift För Robotar.'

16. 'Build a Baxter Robot | Rethink Robotics.'

17. 'Montör Lönér'; 'Svetsare Lönér.'

18. 'Beräkna Arbetsgivaravgift.'

19. 'HUMABOT Challenge.'

It did this so well it won the competition! Cooking isn't that hard. A simple robot can follow a recipe and mix different ingredients together in a predefined order.

Most restaurants have a fixed menu every day. It would be easy to replace human chefs with their mechanical counterparts. China has restaurants where the food is prepared by robots and machines²⁰. The food will always be of the same quality and the restaurant can be open around the clock. More common is to have robotic waitresses²¹. In 2014, Lidköping in Sweden opened the first restaurant in Europe with robots that can talk and take orders from customers and deliver food to the tables²². Without having to pay salaries, the restaurant owner can lower the price of the food. Former McDonald's CEO Ed Rensi has said that robots will take over staff jobs at the fast-food empire, because it's cheaper than employing humans²³.

To summarise, we have seen how assembly-line workers and welders have been replaced by industrial robots ever since the 1960s. But, during the past 10 years, we have seen an increase of collaborative robots that can do more agile tasks and work together with humans. Examples include assembling consumer electronics and cooking food at restaurants, tasks that humans traditionally have done for a living. We have also seen how the development of industrial robots is following the same path as the development of computers, and are getting easier to use and easier to reprogram for new tasks.

Transportation sector and driverless cars

If cars were the *coup de grâce* for the horse, driverless cars will dominate our path for the future. Transportation will be transformed in the coming 10 years²⁴. Driverless cars are not just the future, they are the pre-

20. Ward, 'Robot Restaurant: Robots Cook Food and Wait Tables in Harbin'; Hiden, 'Robot Restaurant Where Machines Cook and Serve Food to Customers.'

21. Nguyen, 'Chinese Restaurants Are Replacing Waiters with Robots.'

22. 'Lidköping Öppnar Europas Första Robotkrog.'

23. Haworth, 'Building Robot McDonald's Staff 'Cheaper' than Hiring Workers on Minimum Wage'.

24. Fields, 'Ford's Road to Full Autonomy'; Banker, 'How Will Transportation Change Over The Next 10 Years?'

sent. Self-driving cars are legal in 10 states in US²⁵. Several companies are developing their own driverless cars.

Humans are bad drivers. More than 90% of road accidents are caused by human error²⁶. We text while driving, talk while driving, we have a hard time focusing longer than 20 minutes at a time²⁷. A computer can be super-focused, hour after hour. In addition, human responsiveness is in the magnitude of 200-300 ms²⁸, while a computer can react in a few nanoseconds, over a million times faster. Humans don't have time to make a decision before an accident happens. In addition, humans do not act rationally when stressed. During, or just before, an accident, when we should be acting most rationally, the body produces a lot of cortisol and noradrenalin, which prevents us from making well-founded decisions²⁹.

It is easy to see from tests that algorithms already today makes for better driving than a human driver. The main problem is that humans rarely trust machines to drive. We want a person to drive the airplane or train because it *feels* safe. It's a common feeling that a person makes better decisions, especially if they are in the vehicle and risking their own life³⁰.

In future, we will likely look back on the 20th century and think we were crazy to let humans drive cars over 100 km/h. Biological creatures are not made for such velocities, our brains do not have that capacity. We don't have the ability to make well-informed decisions during an accident. First, we just don't have time to make a decision (just a few milliseconds is faster than our reaction time³¹). Second, we don't have the information needed, for example, we get a binary signal from an indication lamp when driving on a slippery road, while

25. 'Self-Driving Vehicles Legislation.'

26. Olarte, 'Human Error Accounts for 90% of Road Accidents.'

27. **Dukette & Cornish**, *The Essential Twenty: Twenty Components of an Excellent Health Care Team*.

28. 'What Is the Average Human Reaction Time?'

29. Kowalski-Trakofler, Vaught & Scharf, 'Judgment and Decision Making under Stress: An Overview for Emergency Managers.'

30. Mehta et al., 'Consumers' Perceptions About Autopilots and Remote-Controlled Commercial Aircraft.'

31. 'What Is the Average Human Reaction Time?'

an autopilot can have the exact friction coefficient between each wheel and the ground. A computer will have much more information on which to base its decision.

The transportation sector will change dramatically during the next 10 years³². Transportation inside factories are already driverless: paper rolls at paper mills are being transported by AGVs (Automated Guided Vehicle) and driverless trucks move iron ore at mines. But these robots (since driverless cars *are* robots) are not generally known about. Another example is the HHLA Container Terminal Altenwerder in Hamburg, where AGVs and cranes work around the clock without human drivers. It is easy to develop self-driving trucks in a controlled environment: there are no pedestrians and computer models can be made of the surroundings, with electromagnetic wires placed on the ground, just like a lawnmower moves around a garden. A company can save money since salaries don't need to be paid for drivers and the trucks and cranes can operate at all hours.

The next step is to allow self-driving cars in cities, a more challenging task, since we can't forbid people from walking around the city. Cities change and self-driving cars must be able to learn and adapt while they are driving. This is far from impossible.³³ Several campuses in the US have launched self-driving shuttle buses in the past year.³⁴ These have high-definition maps of the surroundings, as well as distance sensors to detect pedestrians.³⁵ Volvo, among others, is developing wireless road trains, where several trucks drive in convoy, with just the first having a human driver, with the others communicating wirelessly³⁶. Uber, Tesla, Google and others are developing self-driving vehicles to replace buses, taxis, trucks and private cars.

The main problem with driverless cars is not of a technological nature, the main problem concerns the legal framework. To prove that it

32. Banker, 'How Will Transportation Change Over The Next 10 Years?'; Fields, 'Ford's Road to Full Autonomy.'

33. 'DARPA Grand Challenge (2007).'

34. Yadron, 'Self-Driving Cars Coming to a College Campus near You as Price of Tech Drops.'

35. Miller, 'Autonomous Cars Will Require a Totally New Kind of Map'; Fields, 'Ford's Road to Full Autonomy.'

36. 'The SARTRE Project.'

isn't that hard to build a driverless car, I built one myself in the spring 2015³⁷. While far from the advanced cars Volvo and Google are developing, it proved possible to build a car that functioned in a city. Other private individuals have done the same; George Hotz built a self-driving car that functioned in traffic.³⁸ The technology is not the limitation here, it is the law that prevents further advances.

The law represents a competitive disadvantage, with companies preferring to develop driverless cars in countries where it is legal. Why would an entrepreneur with a vision to change the transportation sector establish a company in Europe, where all vehicles without a driver are prohibited? There would be a clear advantage to locating headquarters in the US, where it is legal to undertake practical tests. Companies (and future job opportunities) may be hampered by detrimental European laws.

If we look at the labour market in Sweden, over 55.000 people work as truck drivers³⁹. The US has 3.5 million active truck drivers⁴⁰ and is the most common occupation in many states. There are thus over 3.5 million people that risk losing their job in the next decade. They may find new jobs, but the level of complexity of occupations is increasing and requires more educated employees. Not everyone is interested in further education. How can these people cope with this rapid technological change?

To conclude, there are large companies already developing self-driving cars and trucks, and autonomous vehicles will appear on the roads in the coming 10 years. Accidents will fall, since driverless cars have both more time and more information for decision-making. The downside is that many people may lose their jobs when cars can transport goods by themselves. The question is, how will society cope with even more unemployment?

37. Nohrstedt, 'Självstyrande Bil Byggt På En Vecka.'

38. Zelenko, 'On the Road with George Hotz's \$1,000 Self-Driving Car Kit'; Vance, 'Meet the 26-Year-Old Hacker Who Built a Self-Driving Car...in His Garage.'

39. '30 Största Yrkena.'

40. 'Trucking Statistics.'

White-collar workers

White-collar workers (such as administrators and office workers) are more likely to be replaced by immaterial robots than by physical robots. Replacing a white-collar worker does not need mechanical muscle, but instead, software and algorithms. It is obviously cheaper to produce software robots, because this doesn't require actuator and hardware solutions, and once an instance has been developed, it is cheap to reuse the software for more, copying it at almost no cost. In addition, white-collar workers usually have higher salaries than blue-collar workers⁴¹ and therefore a company has a bigger economic motivation to replace them. It may be impossible to replace every task a human worker does completely. Even if only a small fraction of daily work can be automated, the employer will still save money, since it doesn't need to pay a salary for as much time as before.

Automated scripts and tasks have been used for long time, particularly by IT companies. This has been expanding to other companies during the past decade. It can be something as simple as a spam filter in your inbox, so you don't manually have to open every email, to more complicated policies that install computers automatically, in schools and companies, with all the necessary software.

In addition, the development of machine learning during the past 10 years has made it possible to program software that can learn and adapt to new situations⁴². Machine Learning is a research topic within the field of Artificial Intelligence. To create a program that learns by itself neural networks are often used, where the program resembles how the human brain works with synapses and neurons. There are two types of self-learning algorithms: either supervised, where someone needs to tell the algorithm, "This was correct", or unsupervised, where the algorithm itself evaluates if its action was good or bad.

Stock-exchange robots are already taking decisions by themselves. When to buy, when to sell and to whom? No one has taught the robots exactly when and where they should undertake the transactions. They

41. Berg, 'Gap between Blue-Collar and White-Collar Pay Increases.'

42. Kober, Oztop & Peters, 'Reinforcement Learning to Adjust Robot Movements to New Situations.'

analyse the stock market and learn from how the market behaves; they develop and adapt and are improving rapidly. They are an unsupervised artificial intelligence, with the pre-programmed goal to make as big profit as possible. Stock exchanges have been dominated by humans since their inception, but since the beginning of 21st century, computers have started to make automatic trades. This began just as small scripts, becoming more and more advanced. By 2012, algorithms were managing over 85% of total market volume⁴³, and this percentage is increasing.

No programmer can learn what an entire job entails. It would take forever and when eventually finished, the tasks would have changed. Writing a program that replicates someone's work would take a long time and is not sustainable. But what the programmer can do is to develop a self-learning algorithm that can be installed in a computers and observe what an employee does every day. It can analyse every action and keystroke. With this huge amount of information, the algorithm can figure out how to execute any job. Maybe not everything, but enough to free up an employee's time. Fewer people can do more work when computer algorithms are utilised.

Algorithms were by 1997 already beating humans at chess, when IBM's Deep Blue defeated Garry Kasparov. A more recent example is Google's AlphaGo, that played Go and won against Lee Sedol in March 2016. To play chess, Deep Blue analysed all possible movements and selected the ones with the highest possible chance of winning the game. While this constitutes a huge amount of possibilities, they are not impossible for a computer to calculate. The case of Go is different, here there are way too many possibilities for an exhaustive search, so the robot/algorithm needs some form of what we humans call **intuition** or gut feeling. This feeling has been implemented by observing thousands of matches between skilled humans. The next step was then to compete against instances of itself. The same kind of strategy can be used to learn intuition in negotiations and mediation between parties.

Did you read the news this morning? The chance is that the articles

43. Glantz & Kissell, *Multi-Asset Risk Modeling: Techniques for a Global Economy in an Electronic and Algorithmic Trading Era*.

you read were written by a software robot. The Swedish news agency TT is using robots to produce local news⁴⁴. The robot will use statistics and local data to customise articles and notices with a more local touch. The news agency can't write 290 different notices for all municipalities in Sweden and today they can't offer specific notices for a region, municipality or county council. With robots, news agencies can deliver more material to their customers⁴⁵. Also the AP international news agency has been using robots for a long time. Many articles about sport and economics can be written by computers⁴⁶. For example, the software can assess the statistics and history of a team, or a specific player, and write about similarities and differences to previous matches. Today, newspapers can only afford to visit matches in the highest leagues, but demand for news about local teams playing in a lower division can be satisfied by software robots, so that interested readers can access news about their team in the local newspaper or on the internet⁴⁷.

Today, the news produced by software robots is supplementary and adds extra value for customers. Yet, computer programs are getting better at producing text, and there have been experiments with robots that write fiction⁴⁸.

The owners of news agencies will save money if they can use software instead of journalists to produce good-quality material. What will the impact be on the quality of articles and investigative journalism? Complex article series are hard to produce nowadays, given newspapers are dependent on online advertising, which favours 'clickbait' websites, such as Upworthy and BuzzFeed.

To summarise, robots can already make independent decisions (stock-exchange robots), analyse complex events (the Alpha Go robot) and write texts and reports (news robots). Many jobs in these three domains can all be replaced to a degree sufficient that it will impact the labour market.

44. Äng, 'TT Bygger 'reporter-Robot.'

45. 'Mittmedia Börjar Med Robotjournalistik.'

46. Nordström, 'Robotarna Tar Över Journalistiken.'

47. 'Robotar Objektivare Än Journalister.'

48. Gervás, 'Story Generator Algorithms'; Kazemi, 'NaNoGenMo 2015'; Ahlström, 'Det Är En Tidsfråga Innan vi Ser Romaner Skrivna Av Datorer.'

Highly skilled workers

The examples of truck drivers and machine operators centre around replacing low-skilled jobs. But, the creation of digital minds that can make decisions by themselves, mean educated workers may also fear that their jobs will disappear in the future.

If we look at lawyers and law firms, almost all of them have already *hired* robots that do research for them. Automatic computer programs assist them with research and investigation. When one thinks about lawyers, we think they spend most of their time in court giving closing arguments. But, in fact a big portion of their daily work is spent in the office. For example, economic crime lawyers need to read through emails, correspondence and bank transactions looking for deviations. Until the 1980s most of this work was done manually, reading folders of bank transactions and economic reports. It took time, several weeks, to find the one missing receipt or evidence of a bribe. After the arrival of large-scale computers and the internet, all this is done by computers, and the algorithms are getting smarter and smarter every day. Using computers greatly reduced the time needed to find patterns and deviations. Law firms can increase their profits, because they don't need a person working for several weeks when a computer can do the same job in a couple of hours. But saving time or money wouldn't be the main reason to switch to computer algorithms. A human reading through several thousands of letters and bank transactions can lose focus and it is easy to miss a small detail or to make connections between disparate bits of data. A computer is perfect for this, will keep focus and notice small deviations and find patterns⁴⁹.

Analysing huge amounts of data is something humans are not designed for. Our brains don't have the capacity to read and process all the data that is needed for the complex events in our world⁵⁰. We haven't changed notably since the days of the African savannah. We have the same brains now as then. But the world nowadays is much more connected and complex. We are getting bombarded by more during one day than

49. Larsson, "Advokatrobotar" gör Det Billigare Och Lättare Att Få Juridisk Hjälp.

50. Arbesman, 'Is Technology Making the World Indecipherable?'; Badger, "Transit Systems Are Growing Too Complex for the Human Mind."

during our whole life in ancient times. We can't process all this information. But computers can. This research field is called 'big data' analytics.

One robot that is particularly good at big data analysis is IBM's Watson. Watson was designed to play the quiz game Jeopardy! and in 2011 beat former winners Brad Rutter and Ken Jennings.⁵¹ Watson was not allowed to connect to the internet and only used its internal information. But storing information is not a big deal for a computer, the impressive thing was the capability to process human language and understand what it means. In Jeopardy!, competitors are given a subtle and hard-to-understand fact, and must supply the question that would produce the same answer. This is a hard task for a computer. A computer needs to analyse the facts, and in many cases, has to have enough associative ability to understand what the question really is. It can be anything from quotes from a movie to abbreviations of chemical elements.

But Watson's abilities extend beyond Jeopardy! to the medical field. Hospitals in New York and Ohio use Watson to recommend treatment for cancer patients. A human doctor gives Watson information about the patients and Watson responds with a few confidence-scored recommendations. In addition to the information provided by the human doctor, Watson can access research material, clinical studies, journal articles and data about different treatments and drugs. Every day new articles are published in a rate that no human can match⁵². Watson can detect similarities between different kinds of courses of a disease and how different drugs relate to each other with adverse effects. Watson can read through 20 million cancer research papers and come up with the proper diagnosis within ten minutes⁵³. This is far better than any human doctor. Watson is constantly learning and improving through getting results from former patients and reading new medical studies. We may still need human doctors, but the need will decrease when more hospitals use a Watson for diagnosis.

Since Watson is an immaterial robot, there is in theory nothing that

51. 'Watson (Computer).'

52. Van Noorden, 'Global Scientific Output Doubles Every Nine Years : News Blog.'

53. Ng, 'IBM's Watson Gives Proper Diagnosis after Doctors Were Stumped.'

prevents IBM from taking the code and making an app of it. However, the computational power of a mobile phone is currently not sufficient. IBM can instead place Watson in the cloud and let all app users connect to it. If everyone has a doctor in their mobile phone, what is the purpose of a human doctor? The robot will most likely give a correct diagnosis more often than a human doctor, because it has more data to on which to base its decision.

Medical applications have extended beyond the diagnostic medical sphere to surgery. One of the earliest robotic surgical systems was the ZEUS system⁵⁴ and the competing Da Vinci Surgical System⁵⁵. These systems are not autonomous, instead they assist human surgeons. The robot is remotely operated by the human surgeon, taking their movements and shrinking them to minimal, stable movements. A human has bad eyes and can't see much less than 1mm, and can have tremors in movement. The robot acts as a magnifying glass and guides the human to correct locations. This makes the surgery smoother and less invasive. In addition, human surgeons don't need to be in the same room, or even the same country, as the patient. In 2001, surgeons in New York completed the first tele-surgical operation on a patient in Strasbourg, the Lindbergh operation⁵⁶. In the Third World, with low availability of good surgeons, ZEUS has opened the possibility of improving health in the poorest parts in the world.

Research on cognitive and social robots is advancing. In Sweden, the JustoCat, and Japan, the Paro Therapeutic Robot, are both used for dementia patients.⁵⁷ In studies it has been shown that these patients feel better and have lower stress hormone levels if they have an animal to take care of. Unfortunately, this can't be another living being because they may forget to feed or walk it. But a robot doesn't need this. JustoCat is a robot cat and Paro is a robotic seal that reacts and behaves just like a cat, but with mechanical entrails. When people with demen-

54. 'ZEUS Robotic Surgical System.'

55. 'Da Vinci Surgical System.'

56. 'Lindbergh Operation.'

57. Aremyr, 'Kontakten Mellan En Sälrobot (Paro), En Taktill Värmekatt Och Personer Med En Demenssjukdom.'

tia take care of JustoCat or Paro they are happier and less stressed.⁵⁸

To conclude, we have shown that even advanced professions, such as lawyers and doctors, can be replaced by digital counterparts. We are not going to lose all our human doctors during the coming 50 years, but some doctors will lose their jobs to digital solutions but the impact on patients will be positive, since they will get faster and more accurate diagnoses. The same applies to lawyers, the law firms can deliver better and faster results, but individual employees may lose their jobs.

Creative workers

What can humans do better than robots? Many answer that robots can never have feelings and be creative. They may be right, or wrong, depending on how you define creativity and feelings. What if a robot can imitate human feelings and reactions? If a robot reacts in the same way as a human, on the same occasions, does the robot have feelings or is it just imitating those feelings? Where do we draw the line between learned and *real* behaviours?

I have built a robot that painted artwork. A machine that paints the same painting every time is easy to make, but my robot painted different paintings. It won silver at the World Robot Olympiad in South Korea in 2009⁵⁹. Was my robot creative? Some may say yes, others no.

Emily Howell is a robot (a computer program) that composes music and has released two studio albums. The program was developed by David Cope in the 1990s as a project in Artificial Intelligence and music⁶⁰. You just give Emily the genre and duration and it will compose it for you. Because Emily is a software robot, it is capable of composing more pieces in a week than there are seconds during a week. By contrast, a human musician needs breaks to eat and sleep. Emily is easy to

58. Jöranson et al., 'Effects on Symptoms of Agitation and Depression in Persons With Dementia Participating in Robot-Assisted Activity: A Cluster-Randomized Controlled Trial'; Robinson, MacDonald & Broadbent, 'Physiological Effects of a Companion Robot on Blood Pressure of Older People in Residential Care Facility: A Pilot Study.'

59. 'World Robot Olympiad.'

60. Cope, 'Emily Howell.'

duplicate and there can be several Emilys producing music at the same time. Emily learns and is influenced by what it listens to, but never just copies it. This means that every instance will produce slightly different music⁶¹. Would this mean Emily has creativity?

Humans are not overly special, our creativity is just chemical reactions in our brain. There is nothing that prevents us from simulating chemical reactions in a computer. If we increase the simulation size and speed, we could theoretically simulate the whole brain, and arrive at digital creativity that is identical to how human creativity works. Currently, two research projects attempt to imitate the human brain, the EU's Human Brain Project⁶² and the US's BRAIN Initiative⁶³.

Why would we pay a human to make music that Emily can produce for free? Will it be possible to earn a living doing creative tasks in the future, when robots can do the same tasks?

We can, of course, still devote ourselves to music or art, just not for a living, but rather as a way to express ourselves and because we enjoy it. Like the nobility during the Renaissance and Baroque periods, Mozart and Beethoven did not primarily make music for payment. Leonhard Euler and Sir Isaac Newton didn't discover new maths and physics to earn money, but because they were dedicated to these subjects. The privileged classes and skilled artists paid by patronage had time to devote themselves to maths, music or artwork, in contrast to the lower classes that had to work to feed their families. If robots can undertake work for humans in the near future, more time can be spent on intrinsically motivated work, hobbies and leisure time.

Art and other areas of creativity is the last sector robots will likely conquer. But, even if this is several decades in the future, research is so active that we already have robots that can produce fiction, music and art. In summary, even these occupations are not safe in the future.

61. Adams, 'David Cope: 'You Pushed the Button and out Came Hundreds and Thousands of Sonatas.'

62. 'Human Brain Project.'

63. 'BRAIN Initiative.'

Conclusion

The changes we have described will not happen overnight, and some may never happen. But the robots described already exist, they are not science fiction. Society is constantly evolving, but we often don't notice it. One day everyone had a smartphone and our behaviour had changed, but this didn't happen overnight. It is easy to see changes in the past, but not what is happening right now. One often sees development of society as linear, looking at what happened in the past 30 years and imagining the same amount of development will happen the next 30. But what we can see is that the rate of development is increasing and accelerating. What you have experienced today is probably more than a farmer in the 1800s experienced in one year.

Our generation is not any lazier than ones before. We are just as diligent and have an equally high work ethic. But the technical change has surpassed us. Just as the horse was not replaced by cars because it was lazy, it was replaced because cars are a better means of transportation. The main aim of automation is to replace human workforce with digital and mechanical counterparts, not to create new occupations. The digital revolution will sooner or later remove humans from the labour market and make us all unemployable.

A major problem is that the purchasing power of a community decreases when people become unemployed. There are no easy solutions. Maybe each family will in the future own a robot that is working to make money for its owners. But why would a company hire robots when they can buy them cheaper themselves? Another solution is that of a basic income,⁶⁴ a minimum wage to increase individual purchasing power. This may be hard to implement, though several countries are debating it⁶⁵ and testing it on a small scale⁶⁶.

Rather than being cynical and technophobic, this chapter seeks to emphasise that technological change can be something good: it will lead to better living conditions and open up many possibilities. We will

64. Mellqvist, 'Gratis Pengar När Robotar Gör Jobbet.'

65. Stewart and Elgot, 'Jeremy Corbyn to Investigate Idea of Universal Basic Income.'

66. Diaz, 'Basinkomst Ska Testas I Finland'; Kosk, 'En Väg till Ett Nytt Samhälle.'

How may robots affect the labour market in the near future?

have robots that can treat cancer. But, we can't ignore the risks. Similar risks were discussed in the past, by for example, David Ricardo in the 19th century, but the difference here is that the robots are also capable of decision-making.

Society can't solve a problem it is not aware of. This chapter intends to raise awareness of the digital revolution we are facing. We can form our future. What kind of society do you want?

References

'30 Största Yrkena.' *Statistiska Centralbyrån*, 2016. http://www.scb.se/sv_/Hitta-statistik/Statistik-efter-amne/Arbetsmarknad/Sysselsattning-forvarvsarbete-och-arbetstider/Yrkesregistret-med-yrkesstatistik/59064/59071/133973/.

Adams, Tim. 'David Cope: 'You Pushed the Button and out Came Hundreds and Thousands of Sonatas.'" *The Guardian*, July 11, 2010. <https://www.theguardian.com/technology/2010/jul/11/david-cope-computer-composer>.

Ahlström, Kristofer. "Det Är En Tidsfråga Innan vi Ser Romaner Skrivna Av Datorer." *Dagens Nyheter*, April 7, 2015. <http://www.dn.se/kultur-noje/kulturbedatt/det-ar-en-tidsfraga-innan-vi-ser-romaner-skrivna-av-datorer/>.

Arbesman, Samuel. 'Is Technology Making the World Indecipherable?' *Aeon Essays*, January 6, 2014. <https://aeon.co/essays/is-technology-making-the-world-indecipherable>.

Aremyr, Gun. 'Kontakten Mellan En Sälrobot (Paro), En Taktill Värme-katt Och Personer Med En Demenssjukdom,' 2005.

Badger, Emely. 'Transit Systems Are Growing Too Complex for the Human Mind.' *The Washington Post*, March 3, 2016. <https://www.washingtonpost.com/news/wonk/wp/2016/03/03/transit-systems-are-growing-too-complex-for-the-human-mind/>.

Banker, Steve. 'How Will Transportation Change Over The Next 10 Years?' *Forbes*, October 7, 2015. <http://www.forbes.com/sites/steve-banker/2015/10/07/how-will-transportation-change-over-the-next-10-years/#4cc836b43c6e>.

'Baxter (Robot).' *Wikipedia*. Accessed September 25, 2016. [https://en.wikipedia.org/wiki/Baxter_\(robot\)](https://en.wikipedia.org/wiki/Baxter_(robot)).

Berg, Annika. 'Gap between Blue-Collar and White-Collar Pay Increases,' 2005. <http://www.eurofound.europa.eu/observatories/eurwork/articles/gap-between-blue-collar-and-white-collar-pay-increases>.

'Beräkna Arbetsgivaravgift.' *Räkna.net*, 2016. <https://rakna.net/berakna/arbetsgivaravgift/>.

'BRAIN Initiative.' *Wikipedia*. Accessed September 26, 2016. https://en.wikipedia.org/wiki/BRAIN_Initiative.

'Build a Baxter Robot | Rethink Robotics.' Accessed August 22, 2016. <http://www.rethinkrobotics.com/build-a-bot/baxter/>.

Cope, David. 'Emily Howell.' Accessed September 27, 2016. <http://art-sites.ucsc.edu/faculty/cope/Emily-howell.htm>.

'Da Vinci Surgical System.' *Wikipedia*. Accessed September 26, 2016. https://en.wikipedia.org/wiki/Da_Vinci_Surgical_System.

'DARPA Grand Challenge (2007).' *Wikipedia*. Accessed September 25, 2016. [https://en.wikipedia.org/wiki/DARPA_Grand_Challenge_\(2007\)](https://en.wikipedia.org/wiki/DARPA_Grand_Challenge_(2007)).

'Definition of Robot.' *Oxford Dictionaries*. Accessed September 24, 2016. <https://en.oxforddictionaries.com/definition/robot>.

'Demographic History of the United States.' *Wikipedia*. Accessed September 24, 2016. https://en.wikipedia.org/wiki/Demographic_history_of_the_United_States#Historical_population.

Diamandis, Peter H., and Steven Kotler. *Abundance : The Future Is Better than You Think*. Simon and Schuster, 2012.

Diaz, Adolfo. 'Basinkomst Ska Testas I Finland.' *ETC*, September 1, 2016. <http://www.etc.se/utrikes/basinkomst-ska-testas-i-finland>.

Dominiczak, Peter. 'Countryside Population to Increase Dramatically by 2025.' *The Telegraph*, December 30, 2014. <http://www.telegraph.co.uk/news/politics/11318308/Countryside-population-to-increase-dramatically-by-2025.html>.

Dukette, Dianne, and David Cornish. *The Essential Twenty : Twenty Components of an Excellent Health Care Team*. RoseDog Books, 2009.

Fields, Mark. 'Ford's Road to Full Autonomy,' 2016. <https://shift.newco.co/fords-road-to-full-autonomy-36cb9cca330#.rudanbgnj>.

FinWire. ‘EU-Utredning: Arbetsgivaravgift För Robotar.’ *Dagens Industri*, June 22, 2016. <http://www.di.se/artiklar/2016/6/22/eu-utredning-arbetsgivaravgift-for-robotar/>.

Gervás, Pablo. ‘Story Generator Algorithms.’ *The Living Handbook of Narratology*, 2012. http://wikis.sub.uni-hamburg.de/lhn/index.php/Story_Generator_Algorithms.

Glantz, Morton, and Robert Kissell. *Multi-Asset Risk Modeling: Techniques for a Global Economy in an Electronic and Algorithmic Trading Era*. Academic Press, 2013.

Graf, Timothy L. Lead-through robot programming system. US5880956 A, issued 1999. <https://www.google.com/patents/US5880956>.

Haworth, Jessica. ‘Building Robot McDonald’s Staff ‘Cheaper’ than Hiring Workers on Minimum Wage - Mirror Online.’ *The Mirror*, March 26, 2016. <http://www.mirror.co.uk/news/weird-news/building-robot-mcdonalds-staff-cheaper-8044106>.

Hidden, Michael. ‘Robot Restaurant Where Machines Cook and Serve Food to Customers.’ *Mirror Online*, August 21, 2014. <http://www.mirror.co.uk/news/weird-news/robot-restaurant-machines-cook-serve-4084484>.

Hreha, Jason. ‘Is Human Laziness Human Nature?’ *Big Think*. Accessed September 24, 2016. <http://bigthink.com/wikimind/is-human-laziness-human-nature>.

‘HUMABOT Challenge.’ *IEEE-RAS International Conference on Humanoid Robots*, 2014. <http://www.irs.uji.es/humabot/>.

‘Human Brain Project.’ *Wikipedia*. Accessed September 26, 2016. https://en.wikipedia.org/wiki/Human_Brain_Project.

Hunter, David. *International Standard Classification of Occupations*. 1st ed. International Labour Office, 2012. http://www.ilo.org/wcmsp5/groups/public/---dgreports/---dcomm/---publ/documents/publication/wcms_172572.pdf.

‘IRB 14000 YuMi.’ *ABB Robotics*. Accessed September 25, 2016. <http://new.abb.com/products/robotics/industrial-robots/yumi>.

Jøranson, Nina, Ingeborg Pedersen, Anne Marie Mork Rokstad, and Camilla Ihlebæk. ‘Effects on Symptoms of Agitation and Depression in Persons With Dementia Participating in Robot-Assisted Activity: A Cluster-Randomized Controlled Trial.’ *Journal of the American Medical Directors Association* 16, no. 10 (2015): 867–73. doi:10.1016/j.jamda.2015.05.002.

Kazemi, Darius. ‘NaNoGenMo 2015.’ *Github*. Accessed September 27, 2016. <https://github.com/dariusk/NaNoGenMo-2015/issues>.

Kober, Jens, Erhan Oztop, and Jan Peters. ‘Reinforcement Learning to Adjust Robot Movements to New Situations.’ *Robotics: Science and Systems VI* (2011). https://www.google.com/books?hl=sv&lr=&id=q9TxCWAAQBAJ&oi=fnd&pg=PA33&dq=Reinforcement+Learning+to+Adjust+Robot+Movements+to+New+Situations&ots=75i3zmb75a&sig=oQ3C1Iaa4tOjS_6pjiJy_aV6f6U.

Kosk, Mikael. ‘En Väg till Ett Nytt Samhälle.’ *Hufvudstadsbladet*, August 7, 2016. <https://www.hbl.fi/artikel/en-vag-till-ett-nytt-samhalle/>.

Kowalski-Trakofler, Kathleen M, Charles Vaught, and Ted Scharf. ‘Judgment and Decision Making under Stress: An Overview for Emergency Managers.’ *Int. J. of Emergency Management*, 2003, 278–89. <https://www.cdc.gov/niosh/mining/UserFiles/works/pdfs/jadmu.pdf>.

Larsson, Linus. ‘Advokatrobotar’ gör Det Billigare Och Lättare Att Få Juridisk Hjälp.’ *Dagens Nyheter*, August 14, 2016. <http://www.dn.se/nyheter/sverige/advokatrobotar-gor-det-billigare-och-lattare-att-fa-juridisk-hjalp/>.

‘Lidköping Öppnar Europas Första Robotkrog.’ *Robotnyheter*, March 21, 2014. <http://robotnyheter.se/2014/03/21/lidkoping-oppnar-europas-forsta-robotkrog/>.

‘Lindbergh Operation.’ *Wikipedia*. Accessed September 26, 2016. https://en.wikipedia.org/wiki/Lindbergh_operation.

Mar-Nicolle, Diane. 'Wired for Laziness.' *Simon Fraser University*, 2015. <https://www.sfu.ca/sfunews/stories/2015/wired-for-laziness.html>.

McGoogan, Cara. 'App Revenue Will Overtake the Music Industry This Year.' *The Telegraph*, February 10, 2016. <http://www.telegraph.co.uk/technology/2016/02/10/app-revenue-will-overtake-the-music-industry-this-year/>.

Mehta, Rian, Stephen Rice, Scott R. Winter, and Korhan Oyman. 'Consumers' Perceptions About Autopilots and Remote-Controlled Commercial Aircraft.' In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 58:1834–38. SAGE Publications, 2014. doi:10.1177/1541931214581384.

Mellqvist, Gabriel. 'Gratis Pengar När Robotar Gör Jobbet.' *Dagens Industri*, September 20, 2016. <http://www.di.se/artiklar/2016/9/14/gratis-pengar-nar-robotar-gor-jobbet/>.

Miller, Greg. 'Autonomous Cars Will Require a Totally New Kind of Map.' *The Wired*, December 15, 2014. <https://www.wired.com/2014/12/nokia-here-autonomous-car-maps>.

'Mittmedia Börjar Med Robotjournalistik.' *Robotnyheter*, April 24, 2015. <http://robotnyheter.se/2015/04/24/mittmedia-borjar-med-robotjournalistik/>.

'Montör Löner.' *Lönestatistik.se*. Accessed August 22, 2016. <http://www.lonestatistik.se/loner.asp/yrke/Montor-1115>.

Ng, Aldred. 'IBM's Watson Gives Proper Diagnosis after Doctors Were Stumped.' *NY Daily News*, August 7, 2016. <http://www.nydailynews.com/news/world/ibm-watson-proper-diagnosis-doctors-stumped-article-1.2741857>.

Nguyen, Clinton. 'Chinese Restaurants Are Replacing Waiters with Robots.' *Tech Insider*, July 26, 2016. <http://www.techinsider.io/chinese-restaurant-robot-waiters-2016-7>.

Nohrstedt, Linda. 'Självstyrande Bil Byggs På En Vecka.' *Ny Teknik*, May 22, 2015. <http://www.nyteknik.se/innovation/sjalvstyrande-bil-byggs-pa-en-vecka-6393468>.

Nordström, Andreas. 'Robotarna Tar Över Journalistiken.' *Dagens Nyheter* 04, no. 23 (2015). <http://www.dn.se/kultur-noje/robotarna-tar-over-journalistiken/>.

Olarte, Olivia. 'Human Error Accounts for 90% of Road Accidents.' *AlertDriving*, 2011. <http://www.alertdriving.com/home/fleet-alert-magazine/international/human-error-accounts-90-road-accidents>.

Robinson, Hayley, Bruce MacDonald, and Elizabeth Broadbent. 'Physiological Effects of a Companion Robot on Blood Pressure of Older People in Residential Care Facility: A Pilot Study.' *Australasian Journal on Ageing* 34, no. 1 (March 2015): 27–32. doi:10.1111/ajag.12099.

'Robot Operating System.' *Wikipedia*. Accessed September 25, 2016. https://en.wikipedia.org/wiki/Robot_Operating_System.

'Robotar Objektivare Än Journalister.' *Robotnyheter*, March 21, 2014. <http://robotnyheter.se/2014/03/21/robotar-objektivare-an-journalister/>.

'Self-Driving Vehicles Legislation.' *National Conference of State Legislatures*, 2016. <http://www.ncsl.org/research/transportation/autonomous-vehicles-legislation.aspx>.

Stewart, Heather, and Jessica Elgot. 'Jeremy Corbyn to Investigate Idea of Universal Basic Income.' *The Guardian*, September 14, 2016. <https://www.theguardian.com/politics/2016/sep/14/jeremy-corbyn-to-launch-review-into-universal-basic-income>.

'Svetsare Löner.' *Lönestatistik.se*. Accessed August 22, 2016. <http://www.lonestatistik.se/loner.asp/yrke/Svetsare-1117>.

'The SARTRE Project.' Accessed September 25, 2016. <http://www.sartre-project.eu/en/Sidor/default.aspx>.

'Trucking Statistics.' *Truckinfo.net*. Accessed August 22, 2016. <http://www.truckinfo.net/trucking/stats.htm>.

Van Noorden, Richard. 'Global Scientific Output Doubles Every Nine Years : News Blog.' *Nature*, May 7, 2014. <http://blogs.nature.com/news/2014/05/global-scientific-output-doubles-every-nine-years.html>.

Vance, Ashlee. 'Meet the 26-Year-Old Hacker Who Built a Self-Driving Car... in His Garage.' *Bloomberg*, 2015. <https://www.youtube.com/watch?v=KTrgRYazwBI>.

Ward, Alex. 'Robot Restaurant: Robots Cook Food and Wait Tables in Harbin.' *The Daily Mail*, January 13, 2013. <http://www.dailymail.co.uk/news/article-2261767/Robot-Restaurant-Robots-cook-food-wait-tables-Harbin.html>.

'Watson (Computer).' *Wikipedia*. Accessed September 26, 2016. [https://en.wikipedia.org/wiki/Watson_\(computer\)](https://en.wikipedia.org/wiki/Watson_(computer)).

'What Is the Average Human Reaction Time?' *Reference.com*. Accessed September 24, 2016. <https://www.reference.com/health/average-human-reaction-time-64cbcd7617fa4bd2>.

'World Robot Olympiad.' *Wikipedia*, n.d. doi:2016-09-25.

Wyatt, Ian D., and Daniel E. Hecker. 'Occupational Changes during the 20th Century.' *Monthly Labor Review*, 2006. <http://stats.bls.gov/opub/mlr/2006/03/art3full.pdf>.

Yadron, Danny. 'Self-Driving Cars Coming to a College Campus near You as Price of Tech Drops.' *The Guardian*, March 30, 2016. <https://www.theguardian.com/technology/2016/mar/30/self-driving-cars-california-varden-labs>.

Zelenko, Michael. 'On the Road with George Hotz's \$1,000 Self-Driving Car Kit.' *The Verge*, June 6, 2016. <http://www.theverge.com/2016/6/6/11866868/comma-ai-george-hotz-interview-self-driving-cars>.

'ZEUS Robotic Surgical System.' *Wikipedia*. Accessed September 26, 2016. https://en.wikipedia.org/wiki/ZEUS_robotic_surgical_system.

Äng, Anton. 'TT Bygger 'reporter-Robot.'" *Dagens Media*, April 13, 2016. <http://www.dagensmedia.se/medier/digitalt/tt-bygger-reporter-robot-6540498>. [Move this entry to the As?].

Claire Ingram is a PhD Candidate at the Stockholm School of Economics in Stockholm, Sweden. She conducts research into the impact of increased reliance on digital artefacts, especially when it comes to entrepreneurship and the world of work. **Jacob Lindberg** is an Investment Banking Analyst at Nordea, working with hedging of currencies and interest rates. At the time of writing, he worked with professor Robin Teigland as a Research Analyst at Stockholm School of Economics. **Robin Teigland** is Professor in Business Administration, with a specialization in Strategic Information Systems Management, at the Center for Strategy and Competitiveness at the Stockholm School of Economics. Robin's research interests reside at the intersection of strategy, technology, innovation and entrepreneurship.

Building blockchains: In search of a distrib- uted ledger 'stand- ard'?

Introduction

Cryptocurrencies, and most prominently Bitcoin, have received increasing attention in a number of media channels such as *The New York Times* and *The Economist* in the past year. This is quite surprising as Bitcoin only became known to the general public in October 2008 through a white paper published on the internet by a person or group of persons under the pseudonym of Satoshi Nakamoto. The cryptocurrency

was described in the following way: “A purely peer-to-peer version of electronic cash would allow online payments to be sent directly from one party to another without going through a financial institution.”¹ The software code was then released in an open source project in January 2009, and since then it has been further developed and maintained by an open source community of thousands of volunteers distributed across the globe. In mid-2016 Bitcoin had a market capitalisation of around \$10 billion with more than 200,000 daily transactions².

Enthusiasts point to the fact, however, that it is not the cryptocurrency, but its underlying protocol, known as the blockchain, that is making waves in technology circles today. Indeed, some propose that not only does the blockchain enable the digitalisation of trust, but that it may just “drive a productivity revolution across the globe on par with what Henry Ford did with the automobile”³. Cryptocurrencies and blockchain technologies have attracted more than \$1 billion in venture capital investments to date, and while they hold the potential to revolutionise any number of industries, the finance industry has been particularly keen on exploring this potential.

Further development of the technology is currently occurring in a fragmented way, with different stakeholders developing and emphasising different applications of the technology, such as private vs. public blockchains. While Bitcoin is the most widely used cryptocurrency and blockchain infrastructure, hundreds of other distributed ledger technologies using the Bitcoin blockchain, as well as other application-neutral technologies based on the original blockchain idea, are being developed. For example, three commanding significant attention are Ethereum, Hyperledger, and R3 Corda. Ethereum was published as a white paper online by a Russian-born programmer in July 2013 and is based on the Bitcoin blockchain concept, but not the Bitcoin code. It was crowdfunded in 2014 and has been run by a group of core developers as part of the centrally-controlled Ethereum Foundation with a

1. Satoshi Nakamoto, Bitcoin: A Peer-to-Peer Electronic Cash System, vol. 1, 2008.

2. Coinmarketcap.com, July 2016.

3. Paul Brody, Americas Strategy Leader, Technology Sector, Ernst & Young.

market capitalisation of over \$1 billion⁴. Hyperledger was launched by the Linux Foundation in December 2015 and is based on blockchain efforts by a number of organisations, such as Digital Asset Holdings and IBM. Finally, R3CEV is a consortium of more than 50 global financial and other institutions developing Corda, a distributed ledger infrastructure neither built on the Bitcoin blockchain nor with its own native cryptocurrency⁵.

On the surface, it might appear that these distributed ledger technologies are competing directly with one another, with their backing organisations competing in a race to create a technology that will dominate the marketplace. Whether competing or not, in the long run, it is possible that this fragmented development may lead to conflict and wasted resources, thereby undermining the potential of this promising technology. One might even ask whether what we are witnessing is the beginning of a ‘standards war’ in which incompatible technologies compete for market dominance. One of the major benefits of a single standard is that it allows for interoperability, which can be both efficient and make the use of a technology cheaper for users and innovators. However, standards can be hard to manage, and may even lead to either monopolies of control, or sub-optimal solutions as a single dominant standard crowds out other, more-efficient, solutions.⁶ In the blockchain space, detractors disagree both about whether there is a need for a common standard and, if so, what the salient features of such a standard should be.

This chapter examines this issue by taking a closer look at standardisation and standards wars and then linking this to what we are witnessing today with emerging distributed ledger technologies. We begin by discussing two relevant software standardisation cases: UNIX and the war between competing standards, and Java, where firms competed for control, as well as the Linux case where an open source community has steadily driven the emergence of a *de facto* standardisation. We then

4. Ibid.

5. <https://gandal.me/2016/04/05/introducing-r3-corda-a-distributed-ledger-designed-for-financial-services/>, accessed August 2, 2016.

6. Shapiro, C. & Varian, H. R. (1999). ‘The Art of Standards Wars’. *California Management Review*, 41(2), 8-32.

present our understanding of what is occurring today within the distributed ledger technology space and discuss what previous standards and control wars suggest about the development of standards, and control, in the blockchain context. We conclude the chapter by presenting some thoughts and initial policy recommendations as to how things might play out in the future as these technologies develop.

Standards wars and standardisation

As new technologies emerge and penetrate markets, the development of one technology standard tends to lead to significant market benefits, such as: 1) spin-off markets due to the ease of interfacing with the core technology standard; 2) lower business costs as there is only one technology to adopt; 3) improved safety and quality, thereby making products better for consumers; 4) lower environmental impact due to reduced waste, and; 5) ease of comparison which, among other things, gives consumers confidence in the product's quality before purchase.⁷

The path to a technology standard, however, is generally not a straightforward one and in some cases a 'standards war' can even emerge. In their article *The Art of Standards War*,⁸ Shapiro and Varian define standards wars as "battles for market dominance between incompatible technologies", which occur as firms with competing technologies battle to make their technology the industry standard in the hope of increasing their return on investment. Some well-known examples of standards wars include Netscape Navigator vs. Microsoft Explorer in web browsers, Edison vs. Westinghouse in electric power, and RCA vs. CBS in colour TVs.

While standardisation and standards wars have occurred in a number of industries, for the purpose of this chapter we have chosen to focus on the software industry, as it could be argued to be comparable with what we see occurring within the development of distribut-

7. <http://www.iso.org/iso/home/standards/benefitsofstandards.htm> and <http://www.bsigroup.com/en-GB/standards/benefits-of-using-standards/>.

8. Shapiro, C. & Varian, H. R. (1999). 'The Art of Standards Wars'. *California Management Review*, 41(2), 8-32.

ed-ledger technology software. Below we discuss first the UNIX and then the Java path to standardisation. We have chosen these two cases, as they are quite different in terms of where the ‘battle’ occurred. On the one hand, UNIX represents more of a standards war in *content*, i.e., different versions of UNIX software competing against one another to become the technology standard. On the other hand, Java illustrates a standards war in terms of *control*, i.e., how is control exercised over the development of the technology standard. After comparing UNIX and Java, we then turn to the case of Linux and describe its emergent path to standardisation.

UNIX: A ‘standards war’ over content

UNIX is a family of computer-operating systems that has enabled much of the internet’s success due to its underlying philosophy favoring portability and modifiability. For example, both the Linux and Mac OS X operating systems are based on UNIX. UNIX was created at AT&T Bell Labs in the 1970s, and it soon spread to universities, computer companies, and government organisations, such as the Defense Advanced Research Projects Agency (DARPA), which selected UNIX as its operating system in 1979.⁹ However, although UNIX’s tools were powerful, the system used at one organisation could rarely be used at another. For example, any program modifications at, say, Harvard could rarely be shared with computer companies in San Francisco. As a result, numerous different UNIX versions developed, each with its own benefits and drawbacks.

Frustrated by this fragmentation, a group of computer scientists developed the notion of an open system to encourage interoperability, portability, and open software standards. The notion gained momentum and led to the birth of X/Open in 1984, a company whose task it was to define open system environments. Three years later, AT&T merged with Sun Microsystems in order to manage the fragmentation, perhaps with the motive of creating a standard from which they could later profit. The merger worried small software companies, leading them to form

9. https://en.wikipedia.org/wiki/Berkeley_Software_Distribution.

the Open Software Foundation (OSF). In response to OSF, Sun and AT&T then created UNIX International, leading to a battle between OSF and UNIX International that has been called the ‘UNIX Wars’. In the race to win the war through developing competing features, overall UNIX operating performance sacrificed¹⁰.

X/Open, however, remained a neutral player and worked continuously towards standardisation. For example, it created a ‘stamp of approval’ and worked on standardising application-programming interfaces (APIs). These stamps of approval were successful, and in just a few years around \$20 billion worth of software had the X/Open stamp, leading X/Open in essence to become a standardisation authority.

At the same time, this battle enabled the competitor Microsoft to gain market share with its competing operating system product. Thus, in an attempt to overcome fragmentation and create a unified operating system standard for UNIX, the majority of members in OSF and UNIX International – still enemies – formed the Common Open Software Environment (COSE) in 1993. They reasoned that this would benefit the group of UNIX firms as a whole, especially in their competition with Microsoft products. In the same year AT&T sold its UNIX business unit to Novell, which handed over the UNIX trademark to X/Open and then sold the business to the Santa Cruz Operation in 1995. The UNIX wars were winding down and were considered to come to a conclusion in 1996 when the Open Software Foundation (OSF) merged with X/Open. Today, the Open Group is a global consortium of more than 500 member organisations. It functions as the certifying body of the UNIX trademark and publishes the Single UNIX Specification technical standard (SUS), which is the collective name for a family of standards for operating systems¹¹.

10. Miller, B.P., Fredriksen, L. & So, B., (1990). ‘An empirical study of the reliability of UNIX utilities’. *Communications of the ACM*, 33(12), p.32-44.

11. <http://www.opengroup.org/aboutus/>.

Java: A ‘standards war’ for control

While UNIX is an example of competing technologies transcending into open warfare to the detriment of both consumers and the project itself, Java is an example of a top-down path to standardisation in which the ‘war’ was waged over control over the standard development. Java began in 1990 when three software engineers at Sun became frustrated with the C and C++ programming languages and formed a project to develop a new programming language. In 1993 the evolution of the internet encouraged the team, now expanded by Sun’s management, to shift its focus to building a small application that could be run inside a web browser. Given their frustration with C++, the team developed Java, a flexible programming language for the then-emergent internet¹². Sun presented Java at the SunWorld conference in May 1995, and the firm announced that the technology would be incorporated into the world’s most popular browser Netscape Navigator.

Sun pursued an open-systems strategy for Java’s development, allowing anyone to download Java for free and to modify and further develop it. In order to gain momentum and attract users, Sun proclaimed Java to be the future due to the superiority of the technology, arguing that it was a “complete networking platform”¹³.

Income from Java was generated via licensing fees and royalties based on sales of Java-related products by the licensee. While the licensee was free to modify the code, it could do so only if it shared its modifications with Sun and thereby with all other licensees, including competitors. In this way, Sun controlled Java, but allowed creativity among other firms. UNIX-style fragmentation was something Sun wanted to avoid.

Wanting to stay on the top of the software pyramid, Microsoft criticised Java for years, refusing to buy a license for the language. However,

12. Characteristics encouraging flexibility included the following: (1) be simple; (2) have a standard sets of APIs; (3) remove concepts that required manipulation of memory to make the language safe; (4) be platform independent; (5) be able to ‘Write Once Run Anywhere’, i.e., not having to rewrite the code; (6) be embeddable in web browsers; and (7) have the ability for a single program to multi-task. https://en.wikibooks.org/wiki/Java_Programming.

13. Garud, R., Jain, S., & Kumaraswamy, A. (2002). ‘Institutional Entrepreneurship in the Sponsorship of Common Technological Standards: the Case of Sun Microsystems and Java’. *Academy of Management Journal*, 45(1), 196–214.

it changed its tune when it fell behind in the browser war of its Internet Explorer against the Netscape Navigator, due in large part to its unwillingness to use Java. This was because web content developers opted to conduct development using Java, which Internet Explorer did not, at the time, support¹⁴. Microsoft later bought a license and proceeded to significantly change the code. Sun, afraid of fragmentation, sued Microsoft with the court subsequently ruling in Sun's favour.

Even though Sun won the battle, suing Microsoft might have hurt them in the long run as other firms in the industry became worried about Sun's protectionist and controlling behavior. Microsoft, together with Intel and Compaq, suggested that Sun allow an international standards body to monitor Java. However, Sun refused this suggestion. Industry concern further increased when Sun began to introduce its own Java products, putting Sun in direct competition with its licensees, who rightly believed that Sun had an unfair advantage over them in the market. Netscape and Novell pressured Sun to make Java open source while HP created a Java clone that Microsoft endorsed. Finally, Sun applied to make Java a standard, but when the standardisation bodies ISO and ECMA demanded neutrality, Sun withdrew its application, further harming Sun's reputation as it was deemed by industry to be proof of Sun's desire for total control over the software.

In reaction to industry concern around excessive control, Sun made some significant changes in 1998. It separated its standardisation attempts from its product development, which were previously within one entity, and it made parts of the source code publicly available. Additionally, it loosened its licensing rules by requiring an organisation only to pass a compatibility test and to no longer be required to share any modifications. In November 2006 Sun continued by making the code for Java open source freely available under the GNU GPL¹⁵. In October 2009 Oracle acquired Sun and described themselves as stewards of Java with a commitment to transparency. Java has been one of the

14. Egyedi, Tineke M. 'Why Java™ was-not-standardized twice'. *Computer Standards & Interfaces*, 23, no. 4 (2001): 253-265.

15. Copyleft is a general method for making a program (or other work) free, and requiring all modified and extended versions of the program to be free as well. <https://www.gnu.org/copyleft/>.

most popular programming languages in the world since 2009, and today it is supported by Linux and available on all major operating systems.

Comparing the UNIX and Java paths to standardisation

Both the UNIX and Java cases were the result of programmers being frustrated with the incumbent technology and wanting to develop new software addressing the gaps. While only the UNIX case was a standards war according to the Shapiro and Varian definition, the Java case highlights the importance of not just a standard, but also who controls the standard. As such, both cases provide some interesting insights into the paths to standardisation that a technology can take.

Of particular interest is that in both cases, openness was an important element of resolving the standardisation *impasse*. However, the routes to resolution differed. UNIX was closed source and in the initial years of its development many different versions of the software, with limited compatibility, emerged that could not be used between organisations. This was the result of a decentralised organisational setup: no single organisation controlled the software’s development, leading to the inability of a standard to be developed and thus inefficiency in the market. Over time, however, the industry did come together with the purpose of creating a unified operating system standard with the X/Open and COSE initiatives and X/Open emerging as a *de facto* standardisation body.

Java’s development, on the other hand, was more efficient in its initial phase as it was under the centralised control of Sun, which steered its development by enforcing strict rules on how modifications were made. However, over time Sun had to lessen its control and find a balance between controlling too much and too little, while also making the software open source due to backlashes from industry. Of interest is that in neither case did regulation play a role as the firms managed to self-regulate themselves. When Sun was too close to a monopoly, smaller firms attacked it. When UNIX became too fragmented, alliances were formed to combat it. Even though monopolies and fragmented

technology could have been seen as market failures, and thereby a potential need for the government to intervene, regulation did not occur, yet both technologies in the end became their own form of standard. We summarise these cases in Table 1.

Table 1: A comparison of Java and UNIX paths to technology standards

	UNIX	Java
Governance	Decentralised, emergent control	Top-down, centralised control
Software development	Inefficient as individuals and organisations developed their own customised versions that leading to many UNIX versions competing against one another	Efficient as all modifications by licensees were to be openly shared with Sun deciding on which modifications to be adopted into the next release of software
Standards development	Too little 'control' over software development led to standards bodies emerging to facilitate the software to gain standards status.	Too much control over software development by Sun led to backlash from industry and over time Sun had to move towards a more open development policy to ensure that its software retained its standards status.

The emergence of Linux as a 'de facto standard'?

While UNIX is protected under a trademark and was primarily developed and distributed in a closed environment by commercial and non-profit organisations for use by large organisations, such as multinationals and universities, the Linux operating system has emerged since

1991 as a strong competitor despite it being a free and open source operating system. Linus Torvalds, a programmer at the University of Helsinki, based the free operating system on some of the basic principles of UNIX and released the Linux kernel in 1991 to be used primarily for personal computers. Since the kernel’s release, the software has been developed by an extensive open source community across the globe.

Linux contrasts with the UNIX and Java cases above, as it is both an open and a decentralised operating system. The Linux kernel falls under the GNU General Public License, which is based on the principle of copyleft, which means that anyone may run, study, share and/or modify the software, but that any work derived from the Linux kernel, must also fall under the same copyleft license. One of the results of the copyleft license is that smaller independent developers tend to focus on the interoperability of their projects such that their projects can be collected and then redistributed in larger scale projects for others to use. This, however, has led to hundreds of various Linux-based operating systems known as distributions or ‘distros’, with Ubuntu being the one that is viewed as the closest to a standard. While the Linux projects generally follow the various software standards, e.g., SUS, POSIX, ISO, they are, however, not certified.

Despite the software being originally developed in an emergent, decentralised environment outside the control of a corporation or formal organisation, Linux has grown to become one of the most widely used operating systems across the globe. While issues such as user-friendliness and availability of programs are considered to be the drawbacks of Linux, most would agree that Linux comes out ahead of Windows on matters of security. The argument is that since the Linux software, like other open source software, is maintained by an extensive network of individuals across the globe, it is continually under scrutiny by people looking for flaws. Indeed, the ‘Linus Law’, named after Torvalds, states that “given enough eyeballs, all bugs are shallow”¹⁶.

By 2000, most computer companies supported Linux, and in 2001 the Linux Foundation initiated the Linux Standard Base (LSB) in order

16. Raymond, Eric S. (1999). *The Cathedral and the Bazaar*. O’Reilly Media. pp. 30.

“to develop, through consensus, a set of standards that will increase compatibility among Linux distributions and enable conforming products to work with any compliant system – in other words, to provide a single target for vendors building products for the Linux platform. In addition, the LSB helps coordinate efforts to recruit vendors to develop compliant products.”¹⁷

Today large organisations, such as NASA, DELL, IBM, and HP, have switched to Linux software while application areas have grown to include smartphones, tablet computers, cars and home appliances¹⁸. Despite the copyleft license and emergent organisational nature of Linux, a recent analysis revealed that 75% of the Linux kernel code developed between December 2008 and January 2010 was written by programmers working for commercial operations such as Dell, IBM, HP, Oracle, and Nokia. These organisations and others basing their operations on Linux have developed business models in which they charge for support, such as installation and maintenance, or even for proprietary modules for business customers, instead of charging for the underlying software.

Due to its decentralised and open nature promoting portability and interoperability along, a rapid uptake due to its being free, and an extensive global community continuously maintaining and improving the software, organisations of all sizes have recognised the benefits of the Linux operating system, especially in contrast to the fragmented UNIX wars days. Thus, it would appear that, in many ways, Linux has slowly but steadily been developing into the *de facto* standard.

Three paths to standardisation

In summary, we can compare the three different software programs discussed above on two dimensions: open vs. closed source and decentralised vs. centralised governance, with each of the three falling into a different quadrant. While Java and UNIX were characterised by a form of standards ‘war’, Linux saw signs of emerging as a *de facto* standard in the UNIX community, as seen in Table 2.

17. <https://wiki.linuxfoundation.org/lsb/lsb-charter>.

18. <https://en.wikipedia.org/wiki/Linux>.

Table 2: Standardisation in Java, UNIX and Linux

	Open	Closed
Centralised	War for Control: Java	--
Decentralised	De facto emergence: Linux	War for Content: UNIX

We now turn from our discussion above of the development of standards within the software industry to taking a look at blockchain technologies.

Understanding cryptocurrencies and blockchain technologies

The initial stated intention behind Bitcoin was to create “an electronic payment system based on *cryptographic proof instead of trust*, allowing any two willing parties to transact directly with each other *without the need for a trusted third party*.”¹⁹ As such, some of the initial uses of Bitcoin were making purchases for anything from pizza to real estate to education and transferring funds across the globe. These transactions between peers using Bitcoins were designed to be visible to everyone, and their verification was undertaken in a decentralised manner, as was the development and maintenance of the underlying technology, the blockchain. Thus, the Bitcoin blockchain technology comprises two parts: a) the blockchain, a transaction software protocol or database that uses public-key cryptography, and; b) Bitcoin tokens, a cryptocurrency. The blockchain enables users to conduct transactions through a decentralised, peer-to-peer network with each transaction being verified by specific users, or so-called miners, solving hashing challenges in the peer-to-peer network. Transactions are recorded into blocks, which are then appended to the blockchain. Only when a transaction is included in a block is it deemed as confirmed. The miners are rewarded with Bitcoins and transaction fees in respect of their invested computational power. While all transactions are made publicly available, the identities of

¹⁹ <https://bitcoin.org/bitcoin.pdf>.

those performing the transactions typically are not, i.e., users are pseudonymous in Bitcoin even though all Bitcoins can be traced. Bitcoin has been programmed such that there is a steady increase of Bitcoins being mined into the system with the limit of 21 million Bitcoins being reached during the year 2140 – unless mining power makes significant technological advancements or the underlying protocol is rewritten. As of September 2016, approximately 15.9 million Bitcoins have been mined for a market capitalisation of around \$10 billion.

Distributed-ledger technology, or technology that enables members of a network to hold identical records of transactions, such as enabled by the blockchain, has existed for years. However, what is novel about the Bitcoin blockchain technology is that not only does it remove the role of trusted third-parties, e.g., a Central Bank, to validate entries, safeguard transactions, and preserve a historic transaction record through its decentralised and distributed peer-to-peer network system using public-key cryptography, but that it also rewards the users for their work in maintaining the system, i.e., the miners receive Bitcoins for verifying transactions.

One of the primary use cases for blockchain technology is within the financial sector, due to its decentralised infrastructure. National banking systems have been extensively developed to enable oversight of all transactions in order to prevent money laundering and illegal financial flows, such as to and from terrorist organisations. However, these national systems centralise transactions and have to abide by strict regulations, and as such they are less efficient, often more costly, and more vulnerable to attacks than decentralised blockchain infrastructures. Blockchain proponents argue that the blockchain technology has the potential to digitalise trust not only in the financial industry but in any industry where trusted third parties play significant roles, e.g., the insurance, real estate and legal sectors. Thus, if the Bitcoin blockchain technology truly can digitalise trust, then it has the potential to become a significant disruptor. One long-standing member of the Bitcoin community made the following statement, “Saying that Bitcoin is a curren-

cy is like saying that the Internet is email. Currency is just the first app!"²⁰

Many are exploring the potential uses for the blockchain technology, and today four basic use cases have been suggested: 1) lightweight financial systems, e.g., crowdfunding, loyalty programs, local currencies, P2P trading between asset managers, internal accounting systems; 2) provenance tracking that enables the tracking and movement of items e.g., luxury goods, pharmaceuticals, cosmetics, electronics, and critical items of documentation, e.g., bills of lading, letters of credit; 3) inter-organisational recordkeeping to collectively record and notarise any type of data, e.g., healthcare, legal documents, and; 4) multi-party aggregation for the sharing of databases and Internet of Things²¹.

A 'Jungle' of Blockchains

The potential for numerous use cases and the expectancy of major efficiencies has led to the development of hundreds of blockchain protocol-based systems, but these systems are not all technologically, or ideologically, consistent with each other. Some of the more prominent examples of blockchain-based technologies are Bitcoin, Ethereum, Hyperledger, and R3CEV Corda, and they are summarised in Table 3.

20. <http://startusingbitcoin.com/blog/1-what-is-bitcoin/>.

21. <http://www.coindesk.com/four-genuine-blockchain-use-cases/>.

Table 3: Selected examples of blockchain developments across the globe

	Governance	Promoted use
Bitcoin	Multiple individuals and organisations collaborating in open source community	The original blockchain, developed specifically for transferring and maintaining a ledger of Bitcoin transactions. A number of modules have been developed that rely on this underlying blockchain.
Ethereum	Ethereum Foundation drives development, with controlled outside input	Platform with its own coding language, intended to be flexible enough for use in multiple independent projects, including smart wallets, smart contracts and smart vote counting and collection.
R3CEV Corda	R3 leads consortium of more than 50 member firms	Closed development of Corda technology with limited publicity, but international financial actors are among the key contributors.
Hyperledger	Linux Foundation drives development with multiple firms collaborating	Development of a cross-industry standard for use in digital exchange across any conceivable industry.

Bitcoin

Similar to Linux and as noted above, Bitcoin has been developed and is maintained by an open source community. In addition to the Bitcoin Foundation that was established in September 2012 in the US based on the Linux Foundation with a focus “to foster education, engage in advocacy, increase adoption and encourage development of Bitcoin and blockchain technology worldwide,”²² there are numerous other local and national organisations as well as virtual developer communities supporting Bitcoin.

One of the potential limitations of the Bitcoin technology is its abil-

22. <http://bitcoinfoundation.org/about-us/>.

ity to scale due to its block size limit. This issue is becoming increasingly more relevant as the technology continues to gain traction and the number of use cases expands. As a result, there have been a number of major events in the Bitcoin community focused on altering the underlying protocol, which illustrate how an open source community is governed.

In mid-2015 some members of the Bitcoin community forked the software and created Bitcoin XT in order to increase the block size limit. Their hope was that they could gain enough support from the rest of the community to transition the software to their fork from the original software. However, the community did not back this fork and by the proposed switchover date of January 2016, only 10% of the miners were using the XT protocol, far below the requirement to make the fork the *de facto* standard within Bitcoin²³. This forking has since led the way for several others to fork the software, such as Bitcoin Unlimited, Bitcoin Classic, and BitPay Core. The core idea of each of these is to increase the block limit size; however, none has been successful in gaining traction from the Bitcoin community to date²⁴. In addition to increasing the block size limit, those behind Bitcoin Classic also would like to influence the way decisions regarding the code are made within the Bitcoin community, proposing to establish what they say would be a more democratic decision-making process that would involve input from miners and users through a voting process²⁵.

What is worth mentioning about the Bitcoin blockchain technology as it exists today is that it is a permissionless and open system with explicit tokens. In other words, anyone can view the records of transactions as well as make a transaction in the system, i.e., it is permissionless. Furthermore, it is a system in which Bitcoins, i.e., explicit tokens, are used to reward those verifying the transactions. As a result, the blockchain technology represents a trade-off in which disintermediation, i.e., the displacement of the trusted third party through a permis-

23. <http://www.coindesk.com/scalability-debate-bitcoin-xt-proposal-stalls/>.

24. <https://bitcoinmagazine.com/articles/unlimited-classic-and-bitpay-core-bitcoin-s-new-kids-on-the-blockchain-1452705977>.

25. <https://bitcoinmagazine.com/articles/unlimited-classic-and-bitpay-core-bitcoin-s-new-kids-on-the-blockchain-1452705977>.

sionless peer-to-peer network, is gained at the cost of confidentiality as information regarding all transactions, such as asset value, timing, and involved parties is publicly available for viewing.

Ethereum

Ethereum is a next-generation blockchain technology based on the principles of Bitcoin and was initially proposed as a concept in 2013 by Vitalik Buterin, a Russian-born programmer, and then formally announced in January 2014. In order to crowdfund the development of the project, several joined to create a legal entity, the Ethereum Foundation, a Swiss non-profit organisation, in June 2014. The subsequent crowdfunding sale during July-September 2014 of its cryptocurrency, Ether, raised over \$18 million²⁶. ETH DEV was then created as a non-profit organisation under contract from the Ethereum Foundation to manage the development of Ethereum. Developers were then hired, and the Ethereum Frontier network was launched in July 2015²⁷. A second release of the Ethereum protocol then occurred in March 2016. In the future, many developing and supporting Ethereum would like to transition the software from a proof-of-work transaction validation algorithm to a proof-of-stake algorithm, which would remove the need for costly hardware and high electricity costs. Buterin explained recently, “The dream is to achieve onchain scaling [while] running on nothing other than consumer laptops.”

Ethereum differs from Bitcoin in its underlying code such that it is scalable and that it enables users to program custom ‘smart contracts’ onto its blockchain. A smart contract is a piece of code stored on a blockchain that reads and writes data in the blockchain’s database when programmed blockchain transactions trigger the event²⁸. Smart contracts enable and enforce a contract between parties without the need for a third party to oversee enforcement, and potential uses include the trading of financial instruments, real estate, and intellectual property, encouraging multinationals such as Microsoft, JPMorgan

26. <https://ethereum-homestead.readthedocs.io/en/latest/introduction/history-of-ethereum.html>.

27. <https://blog.ethereum.org/2016/02/09/cut-and-try-building-a-dream/>.

28. <http://www.coindesk.com/three-smart-contract-misconceptions/>.

and Thomson Reuters to develop this technology.

One additional proposed use of Ethereum is the creation of decentralised autonomous organisations, in which the rules and decision-making apparatus of an organisation are coded, thereby creating a structure with decentralised control since the need for documents and people to govern the organisation are eliminated²⁹. On April 30, 2016 the DAO, the name of one such organisation, was launched by Slock.it, a German startup, and any individual could invest in this venture capital organisation during a crowdfunding round. More than \$150 million of ether was raised by the end of May 2016 from the community. However, by June 18, more than 3.6 million ether, or around \$60 million, were stolen from the DAO by someone exploiting Ethereum’s poorly developed code³⁰. The result of this ‘hack’ has led to a very interesting matter of principle related to blockchain technologies. One of the core principles of the original blockchain technology, Bitcoin and then Ethereum, is that the record of transactions is immutable, i.e., once transactions have occurred, the records are ‘set in stone’ and cannot be changed. However, as a result of the theft, the Ethereum Foundation proposed a hard fork of the software such that the stolen funds could be returned to the DAO investors. The Ethereum community voted to approve the hard fork, although there is some discussion as to how well all the Ethereum stakeholders were represented in the vote. Still, some members of the community were against the fork since they were of the opinion that the fork went against the principle and purpose of a truly decentralised, leaderless system and that blockchains should be immutable. Today, the result of the hard fork is that there are two competing Ethereum blockchains and corresponding currencies: Ethereum Classic, based on the unforked protocol, and Ethereum Core or Ethereum One, based on the hard fork of the protocol.

R3CEV Corda

While Ethereum and Bitcoin are permissionless and transparent due to all transactions being available for anyone to view and conduct, several

29. <http://www.coindesk.com/understanding-dao-hack-journalists/>.

30. <http://www.coindesk.com/understanding-dao-hack-journalists/>.

efforts are being made to create private blockchain technologies, similar to how many companies created their own intranets in response to the openness of the internet. These systems require permission to participate in the system, thus preventing outsider from viewing transactions. One such effort is led by R3, a startup firm that launched a consortium with nine financial institutions in September 2015. Since then it has grown to more than 50 financial and other institutions. The consortium is developing Corda, which is a distributed ledger platform for financial agreements between regulated financial institutions. It is inspired by the basic principles of the blockchain, yet it does not have a cryptocurrency as its founders reject the idea that all data should be made to everyone³¹. One of the goals of this project is to create an industry standard and common set of protocols in the financial industry so that interoperability and data interchange among users, applications and systems is enabled. However, there is some discussion as to whether such a consortium puts the goal of standardisation ahead of the need for competition among large players in this industry.

Hyperledger

The final blockchain effort worth mentioning is that led by the Linux Foundation, Hyperledger, an effort announced in December 2015. Today there are more than 80 member organisations, including Accenture, Cisco, Intel, JPMorgan, Mitsubishi UFJ Financial Group (MUFG) and even R3, and a team of programmers from numerous organisations, such as IBM, Digital Asset Holding, and the London Stock Exchange, developing it. The goal of Hyperledger, according to its Executive Director Brian Behlendorf, is to build on experience from previous open source projects such as Linux and Apache and to create an ‘umbrella’ for software developer communities building open source blockchain and related technologies³². The basic idea is to bridle the tribalism that is dividing the various blockchain efforts across the globe and instead unite forces to create modular, open source components and platforms for distributed ledger and smart contract technologies.

31. <https://news.bitcoin.com/r3cev-corda-is-not-building-a-blockchain/>.

32. <https://www.hyperledger.org/author/bbehlendorf>.

In summary, significant efforts are being made on developing various blockchain technologies, yet they are being developed through different means. This high level of emergent development means that both developers and policymakers would be advised to think strategically about the risks and possible areas of tension inherent in blockchain development moving forward. The next section suggests some of the areas that are ripe for policy consideration and argues that development may include not just legal guidelines, but also technical ones.

Where to from here?

Although the Bitcoin blockchain argues that transparency and visibility help it avoid fraud and misbehavior, existing norms and laws in certain industries, notably in finance, restrict the extent to which this kind of transparency can be used, in the name of privacy. This poses a problem when it comes to the development of the technology in these, and other, new industries: how should competing norms be dealt with? Moreover, while the blockchain technology has thus far been developed and maintained by an *ad hoc* and decentralised group of developers, many firms have sought to centralise development and control in the name of business certainty. There has therefore been a move to close off areas of the blockchain in future developments – against the wishes of existing community members – as different firms and consortia develop solutions to deal with their specific problems. Firms in the second wave of blockchain technology have therefore tried to develop the technology in line with their own norms, while both excluding existing stakeholders and operating contrary to existing community norms, and creating competing and incompatible versions of the blockchain. Therefore, similar to the UNIX, Java, and Linux cases above, we can map these efforts onto the closed vs open source and centralised vs decentralised matrix (Table 4).

Table 4. Simple categorisation of blockchain development efforts

	Open	Closed
Centralised	Ethereum	--
Decentralised	Bitcoin, Hyperledger	R3CEV

While there are practical reasons for these positions, fragmentation may impact on development of this promising technology, as will failure by parties to agree on a possible future standard. However, cooperation may be beneficial for all those involved. For instance, while today’s fiat currencies, e.g., Swedish crowns, euros, can be exchanged for one another and SWIFT and IBAN are standards that enable transactions across international borders interfaces and different national banking systems, the different national systems are still so incompatible with one another that third parties are required to bridge the gap between the systems, both in terms of the technology and the transfer of funds. Thus, blockchain technologies could lead to several benefits for the financial industries: **lower costs** through shared technologies and infrastructures, **faster settlement** that leads to lower capital and liquidity needs, **increased transparency**, especially with cross-border payments, and **greater security** through the use of cryptography and transparency³³.

Learning from these examples

The UNIX, Java, and Linux cases described above centre on two key questions around the development of a standard. First, how should a standard come into existence; is a crowd-based, messy development process more desirable than top-down development, and if so, why? Second, who is in control of development matters beyond just the development itself; independence and the dominance of individual actors in a market can undermine trust in a technical or social system, as

33. Citi Research, Jan 2016.

the Java case shows.

The UNIX battle centred on a disagreement between the parties as to whether there should be a UNIX standard and, if so, what it should contain. As such, it was a battle among rivals within the UNIX community to establish themselves and their code’s features as the standard. Control was not an issue at the outset, and its distributed development set a precedent for shared control over the standard. In contrast, there was no public disagreement around Java’s contents as a ‘standard’³⁴. Instead, Java and Sun, as a consequence of their chosen governance model and perceived heavy-handedness faced a backlash from competitors and potential users. Thus, they struggled to promote adoption of their standard, rather than to form the standard in the first place. Their insistence on control (combined with the fact that Sun was not independent and impartial) led to a community backlash. Ultimately, while UNIX started out and remained relatively decentralised in its governance structure, this led to its development being stalled by internal disagreements. In contrast, Java’s centralised development process was more streamlined, but adoption became harder and more contested as a consequence. While in the Linux case, the open source community developed the software steadily with a more bottom-up approach.

The original Bitcoin blockchain development resembles the UNIX movement’s development in some ways: although parts of UNIX’s underlying code was originally proprietary (some parts were later made open source), competing versions quickly emerged. UNIX required consensus as to the contents of the software before it could be widely adopted as a standard. With Bitcoin, this lack of consensus-building led to fragmentation in the community and has spawned several blockchain technologies. However, Bitcoin continues to be the leader, as measured in terms of market capitalisation and the community can be said to have technical consensus as people continue to be active on the Bitcoin blockchain, such that Bitcoin today is more similar to the Linux case. In the second wave of blockchain development, coordinated at-

34. Garud, R., Jain, S., & Kumaraswamy, A. (2002). ‘Institutional Entrepreneurship in the Sponsorship of Common Technological Standards: the Case of Sun Microsystems and Java’. *Academy of Management Journal*, 45(1), 196–214.

tempts to develop standards have taken one of the other approaches to development. R3CEV, while pursuing a closed and centralised development, could potentially move into either a war for control if it moves to open source development as it has suggested. Ethereum, with its recent fork, shows signs of a war for content, while Hyperledger is the most recent development and may experience a war for control as the Hyperledger organisation comprises multiple kinds of organisations, each with their own vision of how the blockchain can be developed and used. These endeavors have led to further development of the code: each has the potential for a showdown within, and even between, these competing technologies for market share.

Table 5: Similarity to UNIX, Java, and Linux standards developments

	Open	Closed
Centralised	War for Control: Hyperledger	R3CEV
Decentralised	De facto emergence: Bitcoin	War for Content: Ethereum

Having discussed the direct implications of the Java and UNIX examples for our understanding of the development of the blockchain, we turn now to discussing some of the policy implications of broader developments using blockchain technology, as well as some considerations when discussing the viability of standards.

Policy considerations and implications

Understanding the nature of software development

One of the significant differences between the development of the blockchain and that of previous software is the scale to which blockchain developers see their tools being applicable. Unlike UNIX, Java, and Linux, which provide largely software and coding languages, blockchain developers offer entire infrastructure for future development.

The suggestion that blockchain technology – and thus developments built upon it – might be as influential as the development of the internet³⁵ gives some sense of the possible impact of the technology. Indeed, the International Monetary Fund suggests that the blockchain is likely to do this through digitising trust, something that until now has been a largely relational undertaking. Given both the scale of the likely development of the blockchain, and its importance as a digital node of trust, it may be that robust development of the technology should be prioritised over its speedy development.

Indeed, open source developments, such as Linux, are known for their competitive nature and robustness: by relying on large numbers of people with different skill sets, and removing hierarchies, individuals continually challenge how a project is progressing. These challenges, while making development slow, ensure that bugs in the software are weeded out and that only the best features of the software survive the scrutiny of the community³⁶.

On the other hand, the number of actors possibly affected by the development of this infrastructure point in favor of a streamlined and top-down approach to standard-setting. Take, for instance, the use of the blockchain in finance. International financial transfers, the payment of salaries, and general business conduct could be conducted through the blockchain. Regulators, however, may be obliged to rely on manual self-declarations by the actors involved. However, if a standard were developed in a clear and coherent way, with the involvement of tax authorities and the like, they would be able to build their own systems to interface with the system and thus monitor it from within. Indeed, the same might be said of international cooperation. As R3 has identified, some applications of the blockchain require that multiple firms across the globe update their existing infrastructure in order to be consistent, and fast enough, to make use of a possible blockchain-based financial system. Doing so requires a coordinated and concerted effort, as well as clear goals and predictable development.

35. <http://www.imf.org/external/pubs/ft/fandd/2016/06/adriano.htm> [accessed August 3, 2016].

36. J Ljungberg, (2000) ‘Open Source Movements as a Model for Organising’, *European Journal of Information Systems* 9, no. 4: 208–16.

Such coordinated attempts require monopolies of power, something that open source communities have in the past been quick to push back against, as in the Java case. Indeed, it was one such monopoly of power that led to the emergence of the Bitcoin blockchain in the first place; against the perceived hegemony of the global financial system and its failures during the financial crisis. This history in general, and specifically among those who are blockchain-savvy, suggests that any centralised attempt to develop a single standard is likely to drive a number of would-be users away from using the standard. It may encourage them to develop alternatives of their own, ultimately leading to a standards war and nullifying previous attempts at centralised development.

These two imperatives, for robustness and predictability, push the discussion around competing standards in different directions. Indeed, thus far we have seen elements reminiscent of UNIX, Java, and Linux. However, as mentioned above, due to the nature of the Bitcoin community it seems unlikely that a fully centralised attempt at developing a standardised infrastructure would be warmly received, or live up to, its potential as a tool for coordination.

Moreover, one might argue that having multiple standards is natural. In the UNIX case the development was bottom-up, as it is currently with the Bitcoin blockchain, and multiple UNIX standards did exist at the same time. Variety is often useful for consumers, since some people might prefer blockchain A whereas other prefer blockchain B, due to the fact that A and B have different characteristics and may be adapted to different conditions. Different blockchains thereby fill different functions, depending on who gives financial support for the blockchain as well as who writes the code. Variety is a common consequence of competition: as multiple providers seek to stand out, they develop specialised solutions and customer-specific approaches. This is particularly the case with bottom-up processes guided by principles and new ideas as they unfold over time.

Thus, when it comes to responding to the promise of the blockchain, there is a strong case for policymakers to get involved at the ground floor of development. Indeed, when we say ‘ground floor’, we mean both at an early stage and in the *technical* and *philosophical* devel-

opment of the technology. This is for two reasons: first, the flexibility of the technology means that considerable amounts of economic activity could ultimately be moved onto the blockchain, including in the areas already discussed. Policy concerns around economic activity, for instance those that take into account privacy or individual freedoms, therefore need to be discussed *while* the technology is being developed, not after. Second, the scalability of the technology is likely to mean global integration at a massive scale, with implications for how, and if, individual countries can respond to the technology. Irrespective of whether countries ultimately decide to regulate areas of blockchain use on national, regional or international levels, international integration on the level of technology will facilitate the exchange of information that will support the policies ultimately decided upon.

Global integration that supports national autonomy

One of the concerns around Bitcoin has been that its decentralised and global reach has facilitated growth in the drug trade³⁷, and that it has supported terrorist financing and money-laundering³⁸. Countries have therefore moved to regulate the use of Bitcoin in the name of stemming these illegal activities, with Australia the first to do so. However, piecemeal single-nation responses to global concerns have limited effect, as individual countries’ attempts to enforce tax laws show. Instead, what is needed to support national autonomy is a co-ordinated international response. At a minimum, such a response would allow for the kind of information exchange that would allow individual countries to ensure that their national laws are being respected, and obeyed. Some relevant areas of law include tax law, employment law, trade in illegal goods (for instance in endangered wildlife or protected commodities, such as uranium).

This kind of global integration points to the need for active involvement from policymakers in how the blockchain is developed. While this

37. <http://www.economist.com/news/international/21702176-drug-trade-moving-street-online-cryptomarkets-forced-compete>.

38. <http://www.smh.com.au/world/australia-to-regulate-bitcoin-under-counterterrorism-finance-laws-20160808-gqnnz2.html>.

development is in its infancy, it may be enough for policymakers just to keep abreast of the latest developments or provide incentives for ethical developers or developers willing to cooperate to undertake significant projects. However, given the stakes, it may soon be the case that authorities need to work together to ensure that emergent blockchain standards are built such that they allow for interaction between third-party providers and state regulations. Indeed, given the extent to which the blockchain allows for automated activity, it would be advisable for state agencies to explore developments of their own, in order to ensure that they aren't left trying to respond to a technology that is vastly more advanced than their own capabilities can respond to.

Rights and responsibilities

Application of the blockchain technology in the various areas already discussed demonstrates how it is likely to become increasingly important as both a way of automating economic activity, and in keeping records of historical activity. For this reason, some of the existing guidelines in various countries need to be brought into the digital age.

There have long been arguments against a *laissez faire* approach to regulation of vital economic sectors, or the protection of human rights. These arguments still hold when it comes to the development of the blockchain in economically and socially important areas. This is because these areas are too important to leave to individual businesses to police. Some areas that are important today when it comes to how most large economies are structured, include top-down laws and guidelines around the following: privacy, consumer protection, and security.

Lawyers elsewhere have argued about the potential, and limitations, of enshrining law into code³⁹. Laws are often general in nature and require interpretation, while code is often more static⁴⁰. However, the incompatibility of laws and code is disappearing even as new economic innovations require more legal oversight and guidance.

39. e.g. Jonathan L Zittrain, 'The Generative Internet', *Harvard Law Review* 119 (2006): 1974–2040.

40. Although recent advances in, for example, neural networks, may mitigate against this argument.

Conclusion

In summary, the formation of a standard is a complex process, and one that requires balancing a large number of factors. In general, our understanding is that the current blockchain ecosystem is facing not only the development of different types of blockchains, e.g., public vs. private, but within each development there is also the potential for competing content standards as well as competing sources of control. Examining the past cases of UNIX, Java, and Linux, we see that these kinds of competition require careful handling, but that self-regulation can occur as such ‘wars’ play out. However, given the stakes involved in blockchain development, it may be of interest for regulators to keep a close watch on blockchain developments and consider setting guidelines upon consultation with those involved.

Regarding the actual development of a single standard, there is no clear argument in favour of one. Although a single standard facilitates cooperation across multiple actors and regions, it may turn out to be worse for some purposes than alternative solutions might have been. Indeed, there is a stronger argument to be made for the emergence of multiple solutions; the issue is then one of providing an efficient way of facilitating cooperation and interoperability between the different solutions. We encourage the exploration and development of these emerging digital infrastructures, and we are optimistic that they can be used to improve existing transaction systems.

References

Adriano, A. & Monroe, H. 2016. 'The Internet of Trust', available online at <http://www.imf.org/external/pubs/ft/fandd/2016/06/adriano.htm>.

Egyedi, T. M. 2001 Why Java™ was-not-standardized twice. *Computer Standards & Interfaces*, 23(4), 253-265.

Garud, R., Jain, S., & Kumaraswamy, A. 2002. Institutional Entrepreneurship in the Sponsorship of Common Technological Standards: the Case of Sun Microsystems and Java. *Academy of Management Journal*, 45(1), 196-214.

Gendal. 2016. 'Introducing R3 Corda™: A distributed ledger designed for financial services', <https://gendal.me/2016/04/05/introducing-r3-corda-a-distributed-ledger-designed-for-financial-services/>.

Greenspan, G. 2016. 'Why Many Smart Contract Use Cases Are Simply Impossible', *Coindesk*, April 17. Available online at <http://www.coindesk.com/three-smart-contract-misconceptions/>.

Greenspan, G. 2016. 'Four Genuine Blockchain Use Cases', *Coindesk*, May 11. Available online at <http://www.coindesk.com/four-genuine-blockchain-use-cases/>.

Ljungberg, J. 2000. Open Source Movements as a Model for Organising. *European Journal of Information Systems*, 9(4), 208-16.

Miller, B.P., Fredriksen, L. & So, B., 1990. An empirical study of the reliability of UNIX utilities. *Communications of the ACM*, 33(12), 32-44.

Nakamoto, S. 2008. 'Bitcoin: A Peer-to-Peer Electronic Cash System', available online at <https://bitcoin.org/bitcoin.pdf>.

Palmer, D. 2016. 'Scalability Debate Continues As Bitcoin XT Proposal Stalls', *Coindesk*, January 11. Available online at <http://www.coindesk.com/scalability-debate-bitcoin-xt-proposal-stalls>.

Raymond, E. S. 1999. *The Cathedral and the Bazaar*. O'Reilly Media.

Shapiro, C., & Varian, H. R. 1999. The Art of Standards Wars. *California Management Review*, 41(2), 8-32.

Siegel, D. 2016. ‘Understanding the DAO Attack’, *Coindesk*, June 25. Available online at <http://www.coindesk.com/understanding-dao-hack-journalists/>.

Topsfield, J. 2016. ‘Australia to regulate bitcoin under counter-terrorism finance laws’, *The Sydney Morning Herald*, August 8. Available online at: <http://www.smh.com.au/world/australia-to-regulate-bitcoin-under-counterterrorism-finance-laws-20160808-gqnn2.html>.

Unknown Author. 2016. ‘Shedding light on the dark web’, *The Economist*, July 16. Available online at: <http://www.economist.com/news/international/21702176-drug-trade-moving-street-online-cryptomarkets-forced-compete>.

Van Wirdum, A. 2016. ‘Unlimited, Classic and ‘BitPay Core’: Bitcoin’s New Kids on the Blockchain. *Bitcoin Magazine*, January 13. Available online at <https://bitcoinmagazine.com/articles/unlimited-classic-and-bitpay-core-bitcoin-s-new-kids-on-the-blockchain-1452705977>.

Zittrain, J.L. 2006. The Generative Internet. *Harvard Law Review* 119, 1974–2040.

Karim Jebari is a researcher at the Institute for Futures Studies. He works on moral philosophy, the ethics of technology and political philosophy. His article "Existential risks: exploring a robust risk reduction strategy" in *Science and Engineering Ethics* discusses problems related to those in this article. **Joakim Lundborg** is a software engineer, currently working at Swedish financial service startup *Wrapp*, specializing in big data and machine learning. He has an academic background in computational linguistics at Stockholm University.

The intelligence explosion revisited

Introduction

Nick Bostrom¹ and other scholars²³⁴⁵ have argued that humans will, absent defeaters, create a *general* artificial intelligence (AI). A general AI differs from narrow AI (the kind that powers search engines and chess programs) by having a set of problem-solving skills that is at least as general as a human's. Moreover, they argue that when such an AI is created, it could be improved very quickly, to the point when it is more intelligent than any existing human. That entity may in turn create other

1. Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies*, 1st edition (Oxford, United Kingdom; New York, NY: Oxford University Press, 2014).

2. Olle Haggstrom, *Here Be Dragons: Science, Technology and the Future of Humanity* (Oxford University Press, 2016).

3. Eliezer Yudkowsky and others, 'Cognitive Biases Potentially Affecting Judgment of Global Risks', in: Bostrom N, Cirkovic MM (eds). *Global Catastrophic Risks* (University Press, 2008), pp. 91–119.

4. Eliezer Yudkowsky, *Intelligence Explosion Microeconomics*, Technical Report 2013-1 (Machine Intelligence Research Institute, September 13).

5. David J. Chalmers, 'The Singularity: A Philosophical Analysis', *Journal of Consciousness Studies*, 17:9-10 (2010), pp. 9–10.

entities like it, with even greater capacities, since this entity would be superior to any existing human at creating and improving AI. This would trigger a loop of recursive improvements that would lead to an AI that far *surpasses* every human and every human organisation across most (or even all) human cognitive skills. They refer to such an entity as an Artificial Superintelligence (ASI). They then proceed to argue that its existence and actions threaten the continued existence of humanity as a whole, an existential risk which they name ‘AI-Xrisk’.

Moreover, these scholars seem to assume that this AI will be an *agent*. By agent AI, we mean an AI that can autonomously direct its activity to achieve its goals in a real-world environment unconstrained by formal rules. An agent AI may be guided by very simple rules. For example, the Roomba is an autonomous vacuum cleaner robot. It can execute its task independently in most rooms, including when humans and other animals are present. It is thus a (very) narrow agent AI.

By contrast, a *tool* AI cannot act independently in a real-world environment. Such AI are typically very useful as decision-making tools. For example, AlphaGo is an AI that beat one of the world’s best players at the board game Go. Yet AlphaGo cannot do anything in a real-world environment. It only knows what to do in the constrained environments of a number of games. Note that a ‘real-world environment’ is not necessarily a three-dimensional space. An AI could be an agent if it can act autonomously while socially manipulating people in chatrooms or other digital environments. These are ‘real-world’ in the sense that their complexity is not restricted by formal rules.

The distinction between tool and agent AI may seem esoteric, but it is not. If the AI-Xrisk claim is correct we need to focus, as Bostrom suggests, on figuring out how to create an AI that is ‘friendly’ (i.e. whose interests align with whatever is important to us). If AI-related existential risk is primarily associated with tool AI, other strategies to reduce that risk ought to be prioritised. For example, a wider dissemination of AI code could be a means to prevent any single person or organisation using a tool AI to cause harm or to concentrate power in the hands of a few individuals. However, in the view proposed by Bostrom et al., such efforts would be very risky, since this would make it more probable that

an intelligence explosion might occur before ‘friendly AI’ is developed. In other words, while we need more resources to explore both agent AI and tool AI risk, some risk reduction strategies pit these two risks against each other.

Proponents of the AI-Xrisk claim argue that when/if an ASI comes to exist, the destiny of mankind would be in its hands, just like the destiny of many lifeforms less intelligent than ourselves are in our hands. If the ASI is benevolent, this outcome may be a good one, Bostrom and others suggest. However, an ASI whose programming is not explicitly directing it to be benevolent, would probably have interests that are at odds with ours. For example, an ASI instructed to maximise the production of paperclips could realise “it would be much better if there were no humans because humans might decide to switch it off” and “human bodies contain a lot of atoms that could be made into paper clips”⁶.

This is the argument we shall discuss, henceforth referred to as the ‘AI-Xrisk claim’. Note that this claim is distinct from the claim that *tool* AI may pose a risk, as weapons or stock markets are put in the hands of systems whose workings are too opaque or complex for us to understand. According to the AI-Xrisk claim, an AI agent poses a great risk.

Let us now consider the components of the AI-Xrisk claim.

1. In the future, (months, years, decades, centuries) an entity with general artificial intelligence (AI) on par with an average human will be created.
2. The difference in intelligence between the average and the most intelligent human is very small. Therefore, assuming continued progress, an AI that will surpass the most intelligent human will be created soon thereafter. Call this AI+.
3. An AI+ will be better than the most intelligent human in improving AI-like entities, and will either improve itself or create an improved version of itself very soon after.
4. The process described in (3) will repeat itself until there exists

6. Nick Bostrom, ‘Ethical Issues in Advanced Artificial Intelligence’, in *Cognitive, Emotive and Ethical Aspects of Decision Making in Humans and in Artificial Intelligence*, ed. by I Smit, Vol. 2 (Int. Institute of Advanced Studies in Systems Research and Cybernetics, 2003).

an artificial superintelligent entity (ASI) that *far surpasses* any human being soon thereafter.

5. It is virtually impossible to control an ASI, and mankind’s destiny will be in its hands.

Our argument proceeds as follows. First, we argue that the AI-Xrisk claim is very implausible, absent a very rapid transition from human-level intelligence (or slightly below) to superintelligence (within days, weeks or months). This rapid transition is in the AI-Xrisk literature referred to as an ‘intelligence explosion’. Second, we will argue that these timescales are, while within the realm of logical possibility, highly unlikely. In other words, while we remain agnostic about the possibility of ASI, we conclude that *even if* it were possible, it is unlikely that it would pose an existential risk *as an agent*. Assuming that tool AI is associated with some degree of existential risk, we suggest that *tool AI* is the greater risk.

Box 1: Definitions

General AI:	An artificial intelligence that can solve a number of different tasks comparable to that of humans, and generalise experience from one class of task to another. Ex: The Terminator in The Terminator film franchise.
Narrow AI:	An artificial intelligence that can only solve a very narrow set of tasks. Ex: Deep Blue, the chess playing program that beat Garry Kasparov, can only play chess.
	Both of these concepts are more points on a scale than distinct categories. They are also relative to what you might think of as general, for example humans have a ‘general AI’, with regards to the tasks we care about, but we certainly cannot perform any intelligent task you could think of.

On the importance of a fast take-off

We can imagine a range of scenarios for AI development, from a slow (decades) emergence of AI to a very fast one (days or weeks). Here, we argue that it will be possible to control an AI unless it emerges rapidly (a ‘fast takeoff’).

Defining ‘intelligence’

First, we need to clarify some concepts. The ordinary use of the term ‘intelligence’ refers to the cognitive abilities of a brain, or *what animals can do with their brains*. When we compare the intelligence of humans by means of IQ tests, this is clearly what is thought of. An IQ test does not measure what a person can accomplish with the help of Google, or how well that person solves problems in a social context. While proper nutrition, adequate training and other beneficial environmental conditions have made a significant contribution to the modern human brain’s ability to perform cognitively, the biology of human brains has not changed much in modern time. However, this ordinary and biocentric use of the term ‘intelligence’ will be of little use in this discussion. We are interested in claims about artificial intelligence and how it relates to human intelligence. To be able to make such comparisons we need to use a definition that can be applied to both humans and artificial systems. An AI does not necessarily have a brain, and a plausible distinction between processes that are ‘internal’ and ‘external’ to an AI-system is not feasible. Therefore, the term ‘intelligence’, when applied to an AI, means something different in the AI-Xrisk literature. Here, ‘intelligence’ means ‘the ability to solve problems’. We will follow this convention, and use this definition consistently. To avoid confusion, we will use the word *techne* to denote this problem-oriented notion. This is different from ‘intelligence’ in the ordinary sense in a number of ways that will be clarified.

First, the *techne* is substrate-neutral. If a human person solves a problem better by solving parts of it with a machine, a calculator for example, then this person’s *techne* has been enhanced by the machine. In other words, whereas in other contexts, a technology only counts as a

cognitive enhancer if it somehow becomes part of a human body (either as a drug or as an implant), a technology enhances *techne* if it improves a human's ability to solve problems, regardless of its location. The idea that human cognition is extended across its environment has been developed by proponents of extended cognition⁷.

Second, *techne* is not person-centric. If a group of people can solve problems that an individual cannot, we can ascribe greater *techne* to this group with regard to these problems. The individuals involved in solving a problem need not be persons either. For example, an ant colony may have greater *techne* than some humans with regards to solving some problems.

Third, *techne* does not assume 'subjective experience'. Therefore, we can ascribe *techne* to an entity without committing ourselves to some specific view on the nature of consciousness.

Fourth, *techne* is distinct from 'thinking', or being able to make symbolic representations of the world, where the symbols have semantic properties. Human organisations, like states or corporations, do not need to make such symbolic representations. Yet these organizations can solve very complicated problems and can therefore be ascribed a great degree of *techne*.

Thus, we can ascribe *techne* to computer code, to companies, to humans and to ant colonies. We can also make claims about the impact of technology, specialisation and cognitive tools on human *techne*. For example, learning algebra may not make a person more intelligent in the usual sense, but it helps that person to solve problems that require algebra. With regards to these problems, learning algebra has increased that person's *techne*.

An ASI is superintelligent (relative to humans) in the problem-solving sense, i.e., with regards to *techne*. An ASI can solve very hard problems far better and faster than any person or any organisation equipped with the best-available technology. This is what makes an ASI an existential threat. An AI that can merely outsmart every person unaided by technology could still be controlled by a group of people whose collec-

7. Andy Clark and David J. Chalmers, 'The Extended Mind', *Analysis*, 58.1 (1998), pp. 7-19.

tive *techne* outperforms the AI, or by a single person whose technology enhanced *techne* is greater than the AI's.

It is worth stressing the importance that we are consistent in the use of this concept, because an ASI is not defined as having 'a lot of *techne*' in *absolute terms*, but rather in *relative terms*. In other words, we are not worried about entities whose *techne* is great. We are worried about entities whose *techne* far surpasses ours (or that of our best organisations).

Imagine that, for some magical reason, human *techne* would explode just as the *techne* of an AI explodes. We would in this scenario get two parallel explosions until hitting some ceiling. In this scenario, there would be super-smart machines, but they would not have more *techne* relative to humans than they are today. In this world, there is no ASI, even when the existing programs have super *techne* relative to how humans were in the past. An AI is only an ASI if it is superintelligent relative to humans (or human organisations) that exist whenever it exists. In other words, what matters from an AI risk perspective is not the *techne* of an AI is relative to humans at some previous point in time, but at the time it exists.

Why the intelligence explosion matters

While the curve of progress of AI research is marred by some 'AI-winter' dips,⁸ the large picture shows continuous progress over the last decades. This has not only brought about increasingly powerful machines, but has also greatly enhanced human *techne* across many domains. Academic work is, for example, greatly facilitated by Google search, Wikipedia and Google Translate. Human organisations can gather information, organise logistical supply chains, coordinate work and carry out complex projects that would have been impossible without narrow AI. In other words, AI has greatly contributed to human *techne*, as defined here.

Would progress in AI research continue in the same pace, it is hardly implausible that there could be an AI that is superintelligent compared to a 2016 human by the end of this century. However, this does not

8. Daniel Crevier, *Ai: The Tumultuous History of the Search for Artificial Intelligence* (New York, NY: Basic Books, 1994).

imply that there would be an ASI. On the contrary, we argue, if future AI research follows current trends, human *techne* will be enhanced by the very same technology, such that 2116 humans will be in a good position to control an AI that is superintelligent compared to a 2016 human.

Would narrow AI be the only technology that could improve human *techne*, then we would be less justified in being confident about humanity's ability to keep up. Fortunately, AI technology is only a small subset of the technological innovations that have extended human *techne*. In the distant past, inventions such as language, writing, cooking and a complex society have improved human *techne*. In the previous century, improvements in healthcare and nutrition, as well as better education and access to information, have vastly improved human capability. We have reason to believe that non-AI technology may continue to make substantial contributions in improving human *techne* in the future. For example, Bostrom and others have discussed the potential of human cognitive enhancement, i.e., medical or genetic interventions that directly enhance cognitive capacity. Furthermore, human-machine interface technology (including, but not restricted to, brain-machine interface technology) has the potential to greatly enhance human *techne*, not necessarily by replacing parts of the brain with superior computers, but by allowing the brain to interact with external computers (and other tools) faster and more intuitively. It is worth noting that the combined computing operations of all human brains are in the order of 10^{26} /second. Yet, as is evident, the inefficiencies in combining these operations (humans could be much better at cooperating) means that this number has little to do with the collective *techne* of humanity. However, the sheer amount of available human computational capacity means that even minor improvements in cooperation (a 'software improvement', metaphorically speaking) could yield massive returns in terms of collective *techne*⁹. Therefore, we have good reason to believe that in a slow or moderate takeoff scenario, human *techne* is likely to be enhanced by technology either faster than, or at least as fast as,

9. Yudkowsky, Intelligence Explosion Microeconomics 2013.

AI¹⁰. Increasing machine *techne* will be part of this improvement. However, Bostrom seems to assume that a given increase in machine *techne* will not affect human *techne*. This seems to imply that Bostrom either assumes a fast take-off (since that would imply that AI technology cannot be absorbed by humans) or that human *techne*, for some reason, is not easy to improve with AI technology.

“At some point in the future, a machine might reach approximate parity with this human baseline (which we take to be fixed- anchored to the year 2014, say, even if the capabilities of human individuals should have increased in the intervening years): this would mark the onset of the takeoff.”¹¹

A less favourable interpretation would be that Bostrom seems to conflate intelligence in the ordinary sense, which will likely remain largely unchanged, and *techne*, which will likely be enhanced as AI and other technologies become available¹². This potential conflation would not be unique to Bostrom. For example, Muehlhauser & Solomon write: “The human brain uses 85-100 billion neurons. This limit is imposed by evolution-produced constraints on brain volume and metabolism. In contrast, a machine intelligence could use scalable computational resources (imagine a “brain” the size of a warehouse)” (p.10). Here, the authors seem to argue that human intelligence is not scalable. This is certainly true for ‘intelligence’ in the ordinary sense. But is not true of *techne* which seems to be what the authors claim to be referring to: “For our purposes, “intelligence” measures an agent’s capacity for efficient cross-domain optimization of the world according to the agent’s preferences.”¹³

In our view, the risk landscape is very different in the fast take-off scenario. Here, progress in AI research makes a huge leap and the transition from AI to ASI is so rapid that no *techne*-enhancing technology

10. Robin Hanson makes a similar observation. Robin Hanson, *The Age of Em: Work, Love and Life When Robots Rule the Earth*, 1st edition (Oxford: Oxford University Press, 2016). pp. 347-49.

11. Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies*, 1st edition (Oxford: Oxford University Press, 2014). pp. 63.

12. Bostrom, *Superintelligence*. pp. 62-63.

13. Luke Muehlhauser and Anna Salamon, ‘Intelligence Explosion: Evidence and Import’, in *Singularity Hypotheses*, ed. by Amnon H. Eden and others, The Frontiers Collection (Springer Berlin Heidelberg, 2012), pp. 15-42 <http://link.springer.com/chapter/10.1007/978-3-642-32560-1_2> [accessed 11 September 2016].

can be absorbed by human society. In such a scenario, it is possible that not even the AI researchers themselves realise that they are dealing with a machine that has the potential to become superintelligent. This is particularly true in the ‘nine hackers in a basement’ version of this scenario, where a small group somehow manages to create an AI without really understanding that they have done so. This is an unlikely, but possible, scenario, according to Bostrom¹⁴. Such a scenario could, we agree, be potentially catastrophic. If, but only if, a fast take-off would be as likely as Bostrom seems to suggest, we would agree with the AI-Xrisk claim.

In conclusion, a slow (and to a lesser degree a moderate) take-off would allow technology to enhance human *techné* such that AI is not likely to cross the human baseline. A fast take-off would, by contrast, be much more dangerous, even for future humans. Thus, the AI-Xrisk claim requires a fast take-off.

On the implausibility of an intelligence explosion

We have argued that agent AI only is a likely risk under the assumption of the plausibility of a fast take-off. Now, we will argue that such a scenario is not plausible for agent AI.

Agents and tools

While the distinction between *tool* AI and *agent* AI was mentioned in the introduction, a further elaboration of this distinction is appropriate. These two concepts are not binary, but refer to extreme values of a spectrum of possibilities. Any AI is somewhere along the spectrum of agency, just as any AI is somewhere along a spectrum of generality. Analogously, to be closer to the agent end of the tool-agent spectrum means being able to act autonomously in a more diverse set of environments, just as being more general means that an AI can solve a more diverse set of problems. In other words, making an AI more agent-like implies ex-

14. Bostrom, *Superintelligence*. pp. 64.

panding the diversity of environments in which it can act autonomously, while being more general implies expanding the diversity of problems that the AI can solve. These concepts are thus related but distinct.

For example, Tesla Motor's Autopilot function is an agent AI, even if it is very limited. This system can keep the distance from a large variety of obstacles on a highway and manoeuvre effectively to avoid collisions. If a future version of this software allows it to drive on city streets, such an upgrade would make it more of an agent AI. By contrast, Google Search is a tool AI, since it can only execute behaviour under the direction of humans in a (more or less) well-defined environment (the web).

The concepts 'tool AI' and 'agent AI' are distinct from the concepts 'narrow AI' (an AI that can only solve a very narrow set of problems, like DeepMind's AlphaGo) and 'general AI' (the kind of AI we have been talking about up to this point). Since a general AI is an AI that can solve a *sufficiently* diverse subset of problems that humans can solve, such an AI is not necessarily an agent, depending on how stringent demands we pose on an AI for it to qualify as 'general'. Being an agent may be a necessary condition for solving some of the problems that humans are typically capable of solving and therefore necessary to attain human levels of generality. However, agency does not seem to necessarily follow from other abilities in AI systems. Rather, agency seems to be a specific set of skills. In other words, improving a narrow AI or combining multiple narrow AI systems will not automatically produce agency in the system. This matters because proponents of the AI-Xrisk claim seem to assume that agency is an emergent property that can suddenly appear in a sufficiently sophisticated system. For example, in *Superintelligence*, Bostrom describes a hypothetical scenario where an AI has the ability to improve its own code, and thus this very same ability. Somehow, this AI also acquires agency, so it starts planning covert actions against its creators¹⁵. While it is tempting to ascribe agency to any sufficiently sophisticated system, research from the field of situated AI (explained later) has shown that agency does not necessarily follow from

15. Bostrom, *Superintelligence*. pp. 95-96.

high *techné*¹⁶. If agency and generality are distinct, then whether or not a specific general AI would be an agent remains an open question. However, creating agency, in particular agency across a significant number of environments, has proven to be very difficult.

The problem of agency

Philosophy of human cognition has in the last decades equated this complex phenomenon with the ability of human brains to make symbolic representations of actual and possible worlds. According a view that has been mainstream among cognitive scientists, intelligent behaviour is explained by the ability of human brains to create accurate models of the world and act on these¹⁷. The assumption that humans and other animals primarily interact with the world by representing it symbolically was perhaps responsible for the difficulty of creating machines capable of agency.

While symbolic representations are arguably an important feature of human cognition, it is also a relatively novel feature that our non-human predecessors developed on top of an already quite sophisticated brain, capable of dynamic interaction with the world. For example, wasps have probably only the most rudimentary abilities to make symbolic representations. Yet wasps can solve a wide range of problems in surprisingly complex environments, including building nests, hunting insects, and avoiding predators. Other cognitive abilities that do not necessarily require rich symbolic representations, such as perception, motion and volition, are in fact very important for being an agent, i.e., for the ability to act autonomously. Being able to symbolically represent the world is neither a necessary nor a sufficient condition for being good at interacting with the world. This is particularly true if we are interested in understanding non-human agents. For example, ant colonies do not make symbolic representations of the world. Yet ant colonies are resilient and successful superorganisms that can plan, react

16. Rodney A. Brooks, 'Elephants Don't Play Chess', *Robotics and Autonomous Systems*, Designing Autonomous Agents, 6.1 (1990), pp. 3–15 <[http://doi.org/10.1016/S0921-8890\(05\)80025-9](http://doi.org/10.1016/S0921-8890(05)80025-9)>.

17. Fodor, Jerry, 'The language of Thought Revisited', Oxford University Press 2010.

and evaluate situations autonomously. By contrast, there is little Google Search could do without humans, even if it can symbolically represent a large amount of information.

Perhaps inspired by the notion of human cognition as a process that primarily represents the world, developers in the early days of AI created top-down, rule-based AI systems¹⁸. These made representations of the world and guide action with a very large set of complex hierarchical rules. We now know that these types of AI (also known as “Good Old Fashioned AI”, GOFAI) while very successful in a wide variety of applications, have severe limitations as agents. Interacting with a complex environment in real time leads to a combinatorial explosion of the number of rules needed. For example, a naïve rational agent that makes a decision, such as choosing two servings of ice cream, by evaluating every possible outcome must make pairwise comparisons of every combination of ice cream taste. For every new taste introduced, the number of evaluations will increase exponentially. Moreover, an agent AI must be able to represent the world with a set of symbols and act on it. Yet as it acts, it also changes the world and must therefore recompute its representations dynamically. To create a general AI agent, which can act autonomously in the real world, traditional AI systems would require overcoming this combinatorial explosion problem.

In real-world GOFAI systems, heuristics are often used to reduce the computational burden, as exemplified by IBM’s famous chess software DeepBlue. However, heuristics are also not generalisable, which means that while they can significantly improve performance when performing specific tasks, they share the limitations of Artificial Neural Networks (see below): non-transparency, and high recalcitrance to radical improvement. Consequently, heuristics are unlikely to extend the agency of GOFAI in any sense that would allow for an intelligent explosion.

The difficulty in creating agents with the GOFAI approach is relevant to our considerations, since this is the approach that most readily allows for rapid progress as a result of “a few critical insights” GOFAI systems

18. See for example: Newell, Allen; Simon, H. A. (1976), ‘Computer Science as Empirical Inquiry: Symbols and Search’, *Communications of the ACM* 19 (3): pp. 113–126, doi:10.1145/360018.360022
 ‘A physical symbol system has the necessary and sufficient means for general intelligent action.’

can quickly improve with the addition of more powerful hardware, and GOFAI code is transparent and can be optimised effectively. For example, the invention of the PageRank (as used by Google in its search engine) and tf-idf algorithms has revolutionised the capability of finding information in unstructured text, and can be applied in many domains. They have not however led to the immediate emergence of agent AI capable of understanding and acting upon the material in the texts. This means that while tool GOFAI could potentially be radically improved by innovations in mathematics and computer science, these improvements will not contribute to the sudden emergence of agent AI. Instead, other approaches seem much more promising for creating agents.

2.3 Situational AI, neural networks and deep learning

Proponents of the AI-Xrisk claim often make reference to animal brains. Surely, if primitive insect brains can overcome the problem of combinatorial complexity, so could engineered AI systems. However, biological brains cannot serve as ‘proof of concept’ for the possibility of general GOFAI agents, since they do not solve the combinatorial problem with this top-down strategy. A large fraction of human action, such as when we interact with our everyday environment, is guided by perception, motion and simple heuristics. These are cognitive processes that require little or no symbolic representation. For example, climbing stairs requires no conscious awareness, which is why we can do this while asleep.

The failures of top-down systems have prompted researchers interested in agent AI, for example AI that controls robots, to develop an approach known as *situated* or *behavioural* AI¹⁹. These systems are very different from traditional AI, but quite similar to primitive biological agents. For example, the autonomous vacuum cleaner Roomba has no internal representation or model of the world which guides its behaviour. Instead, it has a very simple set of rules such as spiral cleaning (spiraling), wall-following and random angle-changing after bumping into

19. Brooks, Rodney (1990), ‘Elephants Don’t Play Chess’ (PDF), *Robotics and Autonomous Systems* 6 (1-2): pp. 3–15.

an object. This means that the Roomba can effectively carry out its task in a variety of environments, without needing specific rules for recognising and interacting with each possible object. This robot's behaviour is quite similar, although much less intelligent²⁰, to that of some primitive insects. Instead of relying on a complete and detailed internal representation of their environment, ants follow a set of simple rules. In conjunction with the behaviour of other ants, this results in behaviour that is surprisingly capable of solving difficult computational problems.

Presently, the most promising direction in artificial intelligence research are models inspired by, or directly simulating, what happens in biological brains. In some ways, Artificial Neural Networks (ANNs) resemble the emergence of higher-level behaviour of ant colonies. Simply put, these AI systems consist of a large number of nodes that each generate a certain output value for a given input value according to very simple rules. ANNs work by aggregating this behaviour, such that it produces a particular output.

ANNs excel at unsupervised and semi-supervised learning tasks, where it is unknown or very hard to build a rule-based model that describes the desired behaviour. To some extent, they are also able to model the spontaneous formation of categories and behaviours, for example with self-organising maps and adaptive resonance theory (ART). This makes the approach very promising as a means of producing intelligent agents. ANNs, and in their latest formulation, deep-learning models, have and will continue to produce impressive results. However, the topic we are interested here is not whether they can produce good results per se, but if a fast take-off scenario is possible or probable using these methods. We believe this is unlikely within this paradigm, because of a number of factors inherent to the approach. These remarks have been repeatedly made in public comments by leading experts in the field of machine

20. In both raw brainpower and *techné*.

learning. See for example Yann Le Cun and Andrew Ng^{21,22}.

First, these agency-oriented AI systems do not process information by creating discrete and transparent mental representations of the world. Whereas GOFAI systems have transparent code that allows for quick gains in performance if a new critical insight on how to solve a problem is arrived at, ANNs are opaque and can only be improved indirectly, by changing how they learn. If there was, for example, a completely new algorithm for solving the travelling salesman problem, a GOFAI that had this as a part of its functioning would be able to get a radical improvement. An ANN that had been trained to approximate solutions for similar problems would not get a boost in this way. For this reason, improving ANNs is similar to a typical engineering task, where networks are trained and tested through trial and error and where improvements are typically local and incremental.

Second, neither ANN nor situated cognition AI is easy to scale up in a brute-force manner in the way that can be done with rule-based approaches. While GOFAI can be made to work in parallel (to the extent the problem is stated in a suitable way) or simply run on a faster processor, ANN and situated AI require more (or better) training data for its *techne* to improve. For example, a rule-based system that works by evaluating as many scenarios as possible, will get more ‘intelligent’ the more of these scenarios it can evaluate, and thus gains directly in performance as hardware power increases. For an ANN, faster hardware reduces training time, but it does not increase the capability of the network as such. Thus, simply adding hardware to existing software does not yield direct increases in *techne* performance in the same way it does with other approaches.

Thirdly, the recent progress of ANNs has been enabled by two recent developments that are unlikely to be repeated. First, improvements in hardware, most notably the use of dedicated hardware in graphic pro-

21. Brian Caulfield, ‘Riding the AI Rocket: Robots Won’t Kill Us, Says Top Artificial Intelligence Researcher’, *The Official NVIDIA Blog*, 2015 <<https://blogs.nvidia.com/blog/2015/03/19/riding-the-ai-rocket-top-artificial-intelligence-researcher-says-robots-wont-kill-us-all/>> [accessed 19 September 2016].

22. ‘What’s Next for Artificial Intelligence’, *Wall Street Journal*, 14 June 2016, section Life <<http://www.wsj.com/articles/whats-next-for-artificial-intelligence-1465827619>> [accessed 19 September 2016].

cessing units, were well suited for the task of simulating artificial neurons. Second, the availability of abundant quality data allowed for rapid learning and competent systems. In the same way as Moore's law has caught up with traditional CPUs (central processing units), the same is now happening to GPUs (graphics processing units). More and more data is being collected, but any further increments here are likely to be incremental rather than revolutionary.

In other words, the low hanging (revolution-enabling) fruits in this domain seem to have been picked and we should expect steady and impressive, but hardly explosive, progress.

Conclusion

In conclusion, we set out to explore some of the concepts and ideas proposed by Bostrom and others in the AI X-risk literature. In particular, we have investigated the claim that agent AI constitutes a major existential risk, and that AI safety should therefore focus on creating 'friendly AI'. We have argued that while human *intelligence* has only marginally improved with modern technology, human *techne* has been radically improved. Moreover, human *techne* is being constantly improved by innovations in technology, and will likely be improved by intelligent machines. Assuming a slow or moderate take-off, human *techne*, at least among the most sophisticated organisations, is likely to improve at least as fast as AI *techne*. Thus, the AI X-risk claim requires a fast take-off.

However, by making the distinction between agent and tool AI more explicit, we argued that such a take-off is implausible for agent AI. The AI systems that have the best track record with regards to autonomous action are also very difficult to improve in the way that Bostrom suggests. This matters because proponents of the AI X-risk claim have been able to deflect criticisms from leading experts in AI research, by appealing to other technological paths to superintelligence than situational AI. If our analysis is correct, that argumentative strategy is not available. To improve agent AI, especially as implemented in ANN or deep learning, is typically an engineering problem, where progress de-

depends on trial and error rather than a small number of critical insights, as is the case in scientific exploration.

However, there is no reason to be complacent if we are correct. Rather, as we have argued, tool AI remains a significant source of existential risk. Such risks have been discussed by Bostrom and others and include the concentration of power, the introduction of too much complexity in global systems or the emergence of totalitarian surveillance states. What these risks have in common is that if they are not possible to address by creating friendly AI systems. Rather, these systems can, like other powerful weapons, empower unfriendly people. To deal with these challenges will require other strategies than the ‘friendly AI’ approach proposed by Bostrom and others.

References

- Bostrom, Nick**, 'Ethical Issues in Advanced Artificial Intelligence', in *Cognitive, Emotive and Ethical Aspects of Decision Making in Humans and in Artificial Intelligence*, ed. by I Smit, Vol. 2 (Int. Institute of Advanced Studies in Systems Research and Cybernetics, 2003), pp. 12–17 <https://www.researchgate.net/publication/229001428_Ethical_Issues_in_Advanced_Artificial_Intelligence> [accessed 11 September 2016]
- Bostrom, Nick**, *Superintelligence: Paths, Dangers, Strategies*, 1 edition (Oxford: Oxford University Press, 2014)
- Brooks, Rodney A.**, 'Elephants Don't Play Chess', *Robotics and Autonomous Systems*, Designing Autonomous Agents, 6 (1990), 3–15 <[http://doi.org/10.1016/S0921-8890\(05\)80025-9](http://doi.org/10.1016/S0921-8890(05)80025-9)>
- Caulfield, Brian**, 'Riding the AI Rocket: Robots Won't Kill Us, Says Top Artificial Intelligence Researcher', *The Official NVIDIA Blog*, 2015 <<https://blogs.nvidia.com/blog/2015/03/19/riding-the-ai-rocket-top-artificial-intelligence-researcher-says-robots-wont-kill-us-all/>> [accessed 19 September 2016]
- Chalmers, David J.**, 'The Singularity: A Philosophical Analysis', *Journal of Consciousness Studies*, 17 (2010), 9–10
- Clark, Andy & David J. Chalmers**, 'The Extended Mind', *Analysis*, 58 (1998), 7–19
- Crevier, Daniel**, *Ai: The Tumultuous History of the Search for Artificial Intelligence* (New York, NY: Basic Books, 1994)
- Fodor, Jerry**, 'The language of Thought Revisited', Oxford University Press 2010
- Haggstrom, Olle**, *Here Be Dragons: Science, Technology and the Future of Humanity* (Oxford University Press, 2016)
- Hanson, Robin**, *The Age of Em: Work, Love and Life When Robots Rule the Earth*, 1 edition (Oxford: Oxford University Press, 2016)

Muehlhauser, Luke, and Anna Salamon, 'Intelligence Explosion: Evidence and Import', in *Singularity Hypotheses*, ed. by Amnon H. Eden and others, The Frontiers Collection (Springer Berlin Heidelberg, 2012), pp. 15–42 <http://link.springer.com/chapter/10.1007/978-3-642-32560-1_2> [accessed 11 September 2016]

'What's Next for Artificial Intelligence', *Wall Street Journal*, 14 June 2016, section Life <<http://www.wsj.com/articles/whats-next-for-artificial-intelligence-1465827619>> [accessed 19 September 2016]

Yudkowsky, Eliezer, *Intelligence Explosion Microeconomics*, Technical Report 2013–1 (Machine Intelligence Research Institute, 2013)

Yudkowsky, Eliezer, In *Global, Catastrophic Risks*, Edited Nick Bostrom, and Milan M. Ćirković, 'Cognitive Biases Potentially Affecting Judgment of Global Risks', in: *Bostrom N, Ćirkovic MM (Eds). Global Catastrophic Risks* (University Press, 2008), pp. 91–119



Karim Jebari is a researcher at the Institute for Futures Studies. He works on moral philosophy, the ethics of technology and political philosophy. His article “Existential risks: exploring a robust risk reduction strategy” in *Science and Engineering Ethics* discusses problems related to those in this article.



Jacob Lindberg is an Investment Banking Analyst at Nordea, working with hedging of currencies and interest rates. At the time of writing, he worked with professor Robin Teigland as a Research Analyst at Stockholm School of Economics.



Joakim Lundborg is a software engineer, currently working at Swedish financial service startup Wrapp, specializing in big data and machine learning. He has an academic background in computational linguistics at Stockholm University.



Fredrik Löfgren works at Linköping University as a roboticist, he also gives lectures about robotics and develops robots in his own company. He has been constructing robots for over 15 years.



Robin Teigland is Professor in Business Administration, with a specialization in Strategic Information Systems Management, at the Center for Strategy and Competitiveness at the Stockholm School of Economics. Robin’s research interests reside at the intersection of strategy, technology, innovation and entrepreneurship.



Karl Wennberg is Professor in Business Administration at the Stockholm School of Economics, deputy director and co-founder of the Institute of Analytical Sociology (IAS) at Linköping University. IAS employs big data to perform detailed empirical research on complex social, political, and cultural processes.

Machines, jobs and equality

Technological change and labour markets in Europe

Rapid technological change in the form of digitalisation, robotisation, electronic payment systems and artificial intelligence, is transforming the production, distribution and consumption of goods and services in Europe. Labour markets are, in turn, radically affected. But how? And how much? Addressing the many economic and societal implications of rapid technological change requires an understanding of the conditions in which it takes place. This, in turn, is essential to the formulation of cogent future policy.

In this volume, a range of experts develop their views on labour markets, productivity, unemployment, redistribution, means of payment and artificial intelligence. They discuss the implications for legislation and regulation. The subjects are of relevance for academics, experts, and political decision-makers alike. The book intends to instigate discussion around these specific issues and probe ideas on how to organise our societies in the face of such rapid technological change.



9 789187 379369 ▶