Cambridge Journal of Regions, Economy and Society 2020, **13,** 25–35
doi:10.1093/cjres/rsz022
Advance Access publication 18 December 2019

The wrong kind of AI? Artificial intelligence and the future of labour demand

Daron Acemoglu[a] and Pascual Restrepo[b]

[a]*MIT, Cambridge, MA, USA, daron@mit.edu*
[b]*Boston University, Boston, MA, USA, pascual@bu.edu*

Received on May 30, 2019; editorial decision on October 16, 2019; accepted on November 18, 2019

Artificial intelligence (AI) is set to influence every aspect of our lives, not least the way production is organised. AI, as a technology platform, can automate tasks previously performed by labour or create new tasks and activities in which humans can be productively employed. Recent technological change has been biased towards automation, with insufficient focus on creating new tasks where labour can be productively employed. The consequences of this choice have been stagnating labour demand, declining labour share in national income, rising inequality and lowering productivity growth. The current tendency is to develop AI in the direction of further automation, but this might mean missing out on the promise of the 'right' kind of AI, with better economic and social outcomes.

Keywords: automation, artificial intelligence, jobs, inequality, innovation, labour demand, productivity, tasks, technology, wages
*JEL Classifications:*J23, J24

Artificial intelligence (AI) is one of the most promising technologies currently being developed and deployed. Broadly speaking, AI refers to the study and development of "intelligent (machine) agents", which are machines, softwares or algorithms that act intelligently by recognising and responding to their environment.[1] There is a lot of excitement, some hype and a fair bit of apprehension about what AI will mean for our security, society and economy. But a critical question has been largely overlooked: are we investing in the "right" type of AI, the kind with the greatest potential for raising productivity and generating broad-based prosperity? We do not have a definitive answer right now—nobody does. But this is the right time to ask this question while we can still shape the direction of AI research and the future of work.

AI as a technology platform

Human (or natural) intelligence comprises several different types of mental activities. These include simple computation, data processing, pattern recognition, prediction, various types of problem solving, judgment, creativity, and communication. Early AI, pioneered in the 1950s by researchers from computer science, psychology and economics, such as Marvin Minsky, Seymour Papert, John McCarthy, Herbert Simon and Allen Newell, sought to develop machine intelligence capable of performing all of these different types of mental activities.[2] The goal was nothing short of creating truly intelligent machines. Herbert Simon and Allen Newell, for example, claimed in 1958 "there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until—in a visible future—the range of problems they can handle will be coextensive with the range to which the human mind has been applied."[3]

These ambitious goals were soon dashed. AI came back into fashion in the 1990s, but with a different and more modest ambition: to replicate and then improve upon human intelligence in pattern recognition and prediction (pre-AI computers were already better than humans in computation and data processing). Many decision problems and activities we routinely engage in can be viewed as examples of pattern recognition and prediction. These include recognising faces (from visual data), recognising speech (from auditory data), recognising abstract patterns in data we are presented with, and making decisions on the basis of past experience and current information. Though there are researchers working on "Artificial General Intelligence", much of the research and almost all commercial applications of AI are in these more modest domains referred to as "Narrow AI"—even if the relevant applications are numerous and varied. The big breakthroughs and the renewed excitement in AI are coming from advances in hardware and algorithms that enable the processing and analysis of vast amounts of unstructured data (for example, speech data that cannot be represented in the usual structured ways, such as in simple, Excel-like databases). Central to this renaissance of AI have been methods of machine learning (which are the statistical techniques that enable computers and algorithms to learn, predict and perform tasks from large amounts of data without being explicitly programmed) and what is called "deep learning" (algorithms that use multi-layered programs, such as neural nets, for improved machine learning, statistical inference and optimisation).

Even if we focus on its narrow version, AI should be thought of as a *technology platform*—there are many ways AI technology can be developed as a commercial or production technology, with widely varying applications. This matters greatly because it implies that the economic and social consequences of AI technologies are not preordained but depend on how we decide to advance and build on this platform. To some degree, this is true of all clusters of technologies, but it is more emphatically so for AI.[4] To see this, contrast it with a related but distinct new technology, robotics. Robotics often makes use of AI and other digital technologies for processing data, but is distinguished from other digital technologies by its focus on interacting with the physical world (moving around, transforming, rearranging or joining objects). Industrial robots are already widespread in many manufacturing industries and in some retail and wholesale establishments. But their economic use is quite specific, and centres on automation of narrow tasks, that is, substituting machines for certain specific activities and functions previously performed by humans.[5]

Implications of technology for work and labour

How do new technologies impact the nature of production and work? Employment and wages of different types of workers? The standard approach, both in popular discussions and in academic writings, presumes that any advance that increases productivity (value added per worker) also tends to raise the demand for labour, and thus employment and wages. Of course, technological progress might lead to job loss in some sectors. But even when that happens, the standard narrative goes, other sectors will expand and contribute to overall employment and wage growth. Moreover, even if technological progress benefits some workers more than others and increases inequality, the standard approach still predicts that it will tend to raise the labour demand for all types of workers.[6]

This view is critically underpinned by the way in which the economic impact of new technology is conceptualised—as enabling labour to become more productive in pretty much all of the activities and tasks that it performs. Yet, this not only lacks descriptive realism (what technology makes labour uniformly more productive in everything?), but may paint an excessively rosy picture of the implications of new technologies. Indeed, in such a world Luddites' concerns about the disruptive and job displacing implications of technology would be misplaced, and they would have smashed all of those machines in vain.

The reality of technological change is rather different. Many new technologies—those we call *automation technologies*—are not intended to increase labour's productivity, but are explicitly aimed at replacing it by substituting cheaper capital (machines) in a range of tasks performed by humans.[7] As a result, automation technologies, by displacing workers from the tasks they were previously performing, always reduce the labour's share in value added. Put differently, these technologies raise productivity by more than wages and employment. They may even reduce overall labour demand (and thus reduce wages, employment or both). Whether they reduce overall labour demand turns on the strength of the *productivity effect* that they create compared with their direct *displacement effect*. The productivity effect is simple to understand: automation technologies typically reduce costs and as costs decline, firms have an incentive to expand output, which increases the demand for labour coming from non-automated tasks. Equally, lower costs for automated products increase the demand for other complementary products, still produced with labour-intensive methods.[8]

A first conclusion from this conceptual framework is therefore that automation technologies always reduce the labour share relative to capital (and other factors), and may or may not reduce overall labour demand. A second conclusion is that whether they reduce overall labour demand depends on the strength of the productivity effect.

This last observation has important implications: contrary to popular claims that the future of labour is threatened by "brilliant" new technologies, the greater danger for labour comes from technology that is not raising productivity sufficiently. In particular, if new automation technologies are not great but just "so-so" (just good enough to be adopted but not so much more productive than the labour they are replacing), there is a double jeopardy for labour—there is a displacement effect, taking passed away from labour, but no powerful productivity gains redressing some of the decline in labour demand generated by the displacement effects.

Is this far-fetched? Not really. We have previously studied the implications of one of the most important automation technologies, industrial robots.[9] Industrial robots are not technologies aimed at increasing labour's productivity but are designed to automate tasks that were previously performed by production workers on the factory floor. The evidence is fairly clear that industries where more industrial robots are introduced experience declines in labour demand (especially for production workers) and sizable falls in their labour share. More importantly, local labour markets more exposed to industrial robots, such as Detroit MI or Defiance OH, have significantly lower employment and wage growth. Furthermore, the declines in wages and employment fell much more heavily on workers from the lower half of the earnings distribution and those with less than a college degree, thus exacerbating inequality. All of this is despite the fact that industry-level data also suggest productivity gains from robots.[10]

Automation in general and robots in particular also increase inequality through two distinct channels. First, by reducing the labour share, automation increases the relative incomes of capital owners who tend to be richer than those relying on labour income. Second, currently automated tasks typically employ low-skill or medium-skill workers, and declines in their employment and wages tend to contribute to inequality. In the case of industrial robots, both of these channels appear to have contributed to greater inequality.[11]

Automation and new tasks in history

Automation is not a recent phenomenon. Many important breakthroughs in the history of technology have centred around automation. Most notably, the spectacular advances in the early stages of the Industrial Revolution in Britain were aimed at automating weaving and spinning, and the focus then shifted to the factory floors of other industries.[12] Other prominent examples of automation are the mechanisation of agriculture and the interchangeable parts system of American manufacturing (for skilled workers).

But if automation tends to reduce the labour share and has mixed effects on labour demand, why did the labour share remain roughly constant and productivity growth go hand-in-hand with commensurate wage growth over the last two centuries? To understand this relationship, we need to recognise different types of technological advances contributing to productivity growth. Historically, as automation technologies were being introduced, other technological advances simultaneously reorganised production, invented new products and created new tasks in which labour had a competitive advantage. These developments generated new activities for labour—tasks in which human labour could be reinstated into the production process—and robustly contributed to productivity growth as new tasks improved the division of labour.[13] The episode of agricultural mechanisation, which started in the second half of the 19th century, vividly illustrates this pattern. Though mechanisation reduced the labour share and employment in agriculture, overall labour demand rose because a range of new tasks were introduced in both manufacturing and services. In fact, this period witnessed not only the rise of clerical occupations but also a range of more specialised blue-collar and white-collar jobs that increased productivity, the demand for labour and the labour share in manufacturing and services.[14] Occupations featuring new tasks have been at the forefront of employment growth in the US economy in the post-war era as well.[15]

This perspective then suggests a different reinterpretation of the history of technology and a different way of thinking about the future of work—as a race between automation and new, labour-intensive tasks. Labour demand has not increased steadily over the last two centuries because of technologies that have made labour more productive in everything. Rather, many new technologies have sought to eliminate labour from tasks in which it previously specialised. All the same, labour has benefited from advances in technology because other technologies have simultaneously enabled the introduction of new labour-intensive tasks. These new tasks have done more than just reinstate labour as a central input into the production process; they have also played a vital role in productivity growth.

Viewed from this perspective, employment and wage growth have been disappointing over the last two decades partly because productivity growth has been weak, and even more importantly because new tasks have failed to materialise.[16] The future of work will be much brighter if we can mobilise more of the technologies that increase labour demand and ensure vigorous productivity growth.

Varieties of AI

This perspective provides a new way of thinking about the economic opportunities and challenges posed by AI. Most AI researchers and economists studying its consequences view it as a way of automating yet more tasks. No doubt, AI has this capability, and most of its applications to date have been of this mould—for example, image recognition, speech recognition, translation, accounting, recommendation systems and customer support. But we do not need to accept this as the primary way that AI can be and indeed ought to be used.

First, if all we do is continue down the path of automation, with no counterbalancing innovations to generate new tasks, the implications for employment, wages and inequality could be depressing. It will not be the end of work anytime soon,[17] but the trend towards lower labour share, greater inequality and anaemic growth in labour demand will continue—with potentially disastrous consequences for income inequality and social cohesion.

Second, as we go deeper and deeper into AI-based automation, we are moving into areas in which human labour is quite good (for example think of image and speech recognition or hand–eye coordination), and machine productivity, at least to start with, is not always impressive, to say the least. Automation technologies aimed at substituting machines for humans in these tasks are thus likely to be of the so-so kind. As a result, we cannot even count on powerful productivity gains to increase our living standards and contribute to labour demand.

But it does not have to be this way. Since AI is not just a narrow set of technologies with specific, pre-determined applications and functionalities but a technology platform, it can be deployed for much more than automation; it can be used to restructure the production process in a way that creates many new, high-productivity tasks for labour. If this type of "reinstating AI" is a possibility, there would be potentially large societal gains both in terms of improved productivity and greater labour demand (which will not only create more inclusive growth but also avoid the social problems spawned by joblessness and wage declines).

Consider a few examples of how AI applications can create new tasks for labour.

- *Education*: Education is one of the areas with the least AI penetration. That may be partly because automation is not an attractive or even feasible option for most of the core tasks in education. But using AI to create new tasks would be a different way of productively deploying this new technology platform. Consider, for example, classroom teaching. This has not changed for over 200 years. A teacher teaches to the whole class, even if he or she or an aide may occasionally engage in one-on-one instruction or provide help for some subset of students. There is evidence, however, suggesting that many students have different "learning styles", and what works for one student may not work for another, and even what works for one student in one subject will not work for him or her in every subject.[18]

At the moment, individualised teaching, targeted and adapted to each student or for small subsets of students, is impossible, and not just because the resources in terms of teacher time and skill are lacking. It is mostly because nobody has (and cannot easily acquire and process the information) to determine a student's optimal learning style in a specific subject or topic. AI can change this. AI software can be designed to collect and process in real-time data about the specific reactions, difficulties and successes students have in different subject areas, especially when taught in different styles, and then make recommendations for improved individualised teaching. The potential improvements in terms of educational productivity could be quite large (we just do not know). Societal benefits could exceed these direct benefits as AI-powered teaching methods may do better in terms of providing students with skills that will be more valued in future labour markets (rather than the more backward-looking curricula and teaching emphasis currently prevailing in schools). Developing and deploying such technologies would increase the demand for human labour in teaching as well—we would need more teachers with diverse skills to do the individualised teaching, even with help from AI software and other technologies.

- *Healthcare*: The situation in healthcare is similar. Though there has been more effort to introduce digital technologies into healthcare, the focus has not been on creating tasks in which humans can be productively employed (in fact, some of the uses of AI, for example, in radiology, are very much in the mould of automation). AI applications that collect and analyse information can significantly empower nurses, technicians and other healthcare providers to offer a wider range of services and more real-time health advice, diagnosis and treatment. The benefits in terms of greater labour demand and productivity are very similar to the education case.

- *Augmented reality:* The third area in which the use of AI can significantly change the production process in a way that may be favourable to labour is through augmented and virtual-reality technologies in manufacturing. Most advanced manufacturing technologies of the last three decades have focussed on automation. But companies such as Amazon and Tesla have discovered that automating all factory-floor and manual tasks is not economically rational because some tasks are still more productively performed by humans. One difficulty facing companies introducing industrial robots, however, is that these new technologies do not necessarily integrate well with humans for at least two reasons. First, most robotics technology is cordoned off from workers because of safety concerns. Second, human work may not mesh with the degree of precision required by robotics technology. Augmented reality technologies—which use interactive interfaces in order to increase the ability of humans to perceive, monitor and control objects—might enable workers to work alongside machines and perform high precision production and integrated design tasks. This will not just help workers keep some of the tasks that might have otherwise been automated; it could also create new tasks in which humans, augmented by digital technology and sensors, can be employed and contribute to productivity.[19]

Notably, the examples of new tasks mentioned above go well beyond what are sometimes emphasised as "enablers" of AI—human tasks involved in training and monitoring new machines as they automate what the rest of us do. This is critical; work in just enabling AI is unlikely to generate sufficient new tasks and demand for human labour to undergird broad-based prosperity.

Why the wrong kind of AI?

If there are potentially productive and profitable uses of AI beyond simple automation, can we count on market forces and innovation by existing companies to take us there? Is there any reason to worry that AI applications with the promise of reinstating human labour will not be exploited and resources will continue to pour instead into the wrong kind of AI?

Economists tend to place great trust in the market's ability to allocate resources in the most efficient way. But most experts recognise that the market's star does not shine as brightly when it comes to innovation. There are several reasons for market failures in innovation in general, as well as some specific reasons that are important in the context of AI.

- Innovation creates externalities—not just the innovator, but also the workers who use the new technology, the firms that adopt it and, most importantly, other firms and researchers building on it in the future will benefit from it. Markets do not do a good job in the presence of such externalities.
- Markets struggle when there are alternative, competing technological paradigms. When one paradigm is ahead of the other, both researchers and companies tend to follow that paradigm, even if an alternative could be more productive. Moreover, in such a situation, once the wrong paradigm pulls ahead, it may be very difficult to reverse this trend

and benefit from the possibilities offered by the alternative paradigm. To the extent that different approaches to AI constitute alternative, competing paradigms, our trust in the market mechanism getting it right should be even lower.[20]

- To correct market failures in innovation, the US government has historically used public–private partnerships to encourage socially beneficial research. It has played an important role in many leading technologies, including the Internet, sensors, pharmaceuticals, biotech and nanotechnology.[21] But more recently, the US government has been more frugal in its support for research and more timid in its determination to steer the direction of technological change. Part of this shift is due to the reduction in resources devoted to government support of innovation and the increasingly dominant role of the private sector in setting the agenda in high-tech areas (can government officials and researchers meaningfully influence the direction of inventive activity in Silicon Valley?). This shift will further discourage research related to future promise (that is not immediately reflected in profitability) and other social objectives (such as reducing carbon emissions or, more relevant to this essay, the creation of employment opportunities for a broad range of workers).
- Innovation does not just respond to economic incentives. Several noneconomic rewards affect what types of technologies attract the attention and imagination of researchers. It is possible that the ecosystem around the most creative clusters in the United States, such as Silicon Valley, excessively rewards automation and pays insufficient attention to other uses of frontier technologies. This may be partly because of the values and interests of leading researchers (consider, for example, the ethos of companies like Tesla that have ceaselessly tried to automate everything). It is also partly because the prevailing business model and vision of the large tech companies, which are the source of most of the resources going into AI, have focussed on automation and removing the (fallible) human element from the production process. This last consideration may have become even more critical as the vast resources of several leading companies are pouring into academia and shaping the teaching and research missions of leading universities. It is no surprise that the best minds in the current generation are gravitating towards computer science, AI and machine learning, but with a heavy emphasis on automation. An ecosystem that is biased would become much more stifling for the direction of technological change when it becomes all-encompassing.
- There are also additional factors that may have distorted choices over what types of AI applications to develop. The first one is that if employment creation has a social value beyond what is in the GDP statistics (for instance, because employed people are happier and become better citizens, or because broad-based growth in labour demand improves income inequality), this social value will be ignored by the market. The second is related to the tax policies adopted in the United States and other Western nations, which subsidise capital and investment while taxing employment. This makes using machines instead of labour more profitable, and these profits encourage not just automation but also automation research. Finally, and complementing these factors, to the extent that firms take into account the cost of labour (the wage rate), which tends to be higher than the social opportunity cost of labour because of imperfections in the labour market, they will have additional incentives for adopting and developing automation technologies beyond what is socially optimal.
- Another set of factors blocking the path of novel AI applications reinstating labour is that these new technologies might need critical complementary inputs that are not

> forthcoming. Take the example of education mentioned above. It is not only that developing AI to create new labour-intensive tasks in education is not viewed the frontier or one of the "cool" areas of research, say compared with facial recognition. It is also that complementary skills and resources to make this type of reinstating AI profitable may be missing completely. Educational applications of AI would necessitate new, more flexible skills from teachers (beyond what is available and what is being invested in now), and they would need additional resources to hire more teachers to work with these new AI technologies (after all, that is the point of the new technology, to create new tasks and additional demand for teachers). In the case of healthcare, limited resources are not the problem (the share of national income devoted to health is continuing to grow), but the requisite complementary changes are likely to be organisational. In fact, highlighting other barriers to the use of new technologies to create new tasks, the way that hospitals, insurance companies and the whole medical profession, as represented by the American Medical Association, is organised is likely to be in the way. If empowering, and increasing the productivity of, nurses and technicians is perceived to reduce the demand for the services of doctors or challenge the current business model of hospitals, it will be strenuously resisted.

All in all, even though we currently lack definitive evidence that research and corporate resources today are being excessively directed towards the "wrong" kind of AI, the market for innovation gives no compelling reason to expect an efficient balance between different types of AI. If at this critical juncture insufficient attention is devoted to inventing and creating demand for, rather than just replacing, labour, that would be the "wrong" kind of AI from the social and economic point of view. Rather than undergirding productivity growth, employment and shared prosperity, rampant automation would contribute to joblessness, anaemic growth and inequality.

Social causes and implications of the wrong kind of AI

Much has been written about the dangers that unregulated AI may pose in the hands of companies or governments intent on monitoring and controlling behaviour or independence of actors wishing to spread disinformation.[22] Without taking away from the importance of these issues, this essay highlights other social aspects of this new set of technologies. We have already emphasised the negative social implications of automation in general and the wrong kind of AI focussing just on automating labour-intensive tasks because they tend to create loss of employment, wage declines or stagnation and greater inequality.

These effects would become even more costly to the extent that loss of employment opportunities, stagnant wages and rising inequality have adverse political implications. These implications could include both mounting popular discontent that can sometimes fan the flames of disruptive populist movements,[23] and growing economic dominance of certain individuals, corporations or segments of the business world, who can then gain disproportionate political influence or even political dominance.[24] These political costs may need to be included in evaluating the broader desirability of different types of AI practices and policies.

The wrong kind of AI does not just have political implications, but its continued dominance may have political causes as well. The wrong kind of AI, primarily focussing on automation, tends to generate benefits for a narrow part of society that is already rich and politically powerful, including highly skilled professionals and companies whose business model is centred on automation and data. If so, the influence of these actors may further propagate the dominant position of this type of AI. For example, corporations that reckon that their own market position and profits will be best served by AI targeted at large-scale automation and would be hurt by new AI technologies creating new tasks, wage growth and opportunities for competing firms may naturally lend their research and political weight towards AI targeting automation. Whether this political channel has had any effect so far and may play more of a role in the future in the path of AI technologies is an interesting and important area for future inquiry.

Conclusion

AI is set to influence every aspect of our lives, not least the way production is organised in modern economies. But we should not assume that, left to its own devices, the right types of AI will be developed and implemented. Though many today worry about the security risks and other unforeseen (often non-economic) consequences of AI, we have argued that there are prima facie reasons for worrying about the wrong kind of AI, from an economic point of view, becoming all the rage and the basis of future technological development. The considerable promise of AI implies that we need to devote care and serious thought to its implications and to the question of how best to develop this promising technology platform—before it is too late.

Endnote

[1] See Russell and Norvig (2009), Neapolitan and Jiang (2018) and Agarwal, Gans and Goldfarb (2018).

[2] See Nilsson (2009) for the history of AI.

[3] Forester (1985, 86).

[4] Bessen et al. (2018) report that many commercial AI startups view their technology as capable of "enhancing human capabilities", while many others recognise that their technologies have a significant automation component.

[5] See Ayres and Miller (1983), Groover et al. (1986), Graetz and Michaels (2015) and Acemoglu and Restrepo (2018b).

[6] See Acemoglu (2002).

[7] This approach is developed in Zeira (1998), Autor, Levy and Murnane (2003), Acemoglu and Autor (2011) and Acemoglu and Restrepo (2018a,b,c, 2019).

[8] One intuition for why the productivity effect, however large, can never restore the labour share back to its pre-automation level is that displaced workers are deployed in non-automated tasks, which runs into diminishing returns (see Acemoglu and Restrepo, 2018a, 2019).

[9] Acemoglu and Restrepo (2018b).

[10] Our conceptual framework implies that the effects of automation technologies on employment and wages will not be uniform across different instances of their adoption or across distinct types of technologies. Robots may create larger displacement effects than other automation technologies, and the same robotic technology in disparate social and economic settings may generate differently sized productivity effects. Indeed, studies focussing on other periods and using different sources of empirical variation find similar declines in labour share from automation, but not always negative effects on employment. See, for example, Graetz and Michaels (2019), who exploited cross-country, cross-industry variation, and Dauth et al. (2019), who use the same strategy as Acemoglu and Restrepo (2018b) but in Germany rather than in the United States.

[11] See Acemoglu and Restrepo (2018b).

[12] See Mantoux (1927) and Mokyr (1992).

[13] Some new technologies also contribute to the productivity of labour directly. Though it is difficult to systematically decompose the contributions of directly "labour-augmenting" technologies and new tasks (broadly construed to include new activities for labour resulting from product innovation and reorganisations), there are two arguments for the importance of new tasks. First, Acemoglu and Restrepo (2019) provide a decomposition suggesting that labour-augmenting technologies have played a relatively minor role in the US economy since 1947. This is both because of empirical reasons

(related to changes in industry labour shares) and also because such technologies impact the labour share only indirectly (working via the elasticity of substitution between capital and labour). Second, the conceptual framework in Acemoglu and Restrepo (2018a, 2019) clarifies that the relative standing of labour cannot be reinstated just by labour-augmenting advances and instead necessitates the creation of new tasks where labour has a comparative advantage relative to capital.

[14] Rasmussen (1982), Olmstead and Rhode (2001) and Acemoglu and Restrepo (2019).

[15] Acemoglu and Restrepo (2018a).

[16] Acemoglu and Restrepo (2019).

[17] See Dreyfus (1992) and Autor (2015).

[18] See Allport (1937), Cassidy (2004), Honey and Mumford (1986) and Ramirez and Casteneda (1974). For recent evidence based on randomised control trials consistent with these ideas, see Muralidharan, Singh and Ganimian (2019).

[19] See Ong and Nee (2013), Kellner (2018) and https://www.ge.com/reports/game-augmented-reality-helping-factory-workers-become-productive/.

[20] See Nelson and Winter (1977), Dosi (1982) and Acemoglu (2012).

[21] Mazzucato (2015).

[22] See, for example, Harari (2018), Lanier (2018), Pasquale (2015) and Zuboff (2019).

[23] See Judis (2016) on the effects of economic hardship and inequality on populism.

[24] See, for example, Acemoglu and Robinson (2012) and Stiglitz (2012) on the political implications of economic inequality.

Acknowledgements

We thank Chris Ackerman, David Autor, Erik Brynjolfsson, Stu Feldman, Mike Piore, Jim Poterba, Hal Varian and two referees and the editors of the journal for useful comments. We gratefully acknowledge financial support from Google, Microsoft, the Sloan Foundation, the Toulouse Network on Information Technology, the Schmidt Sciences Foundation, the Smith Richardson Foundation, IBM, Accenture and the National Science Foundation.

References

Acemoglu, D. (2002) Technical change, inequality, and the labor market, *Journal of Economic Literature*, **40**: 7–72.

Acemoglu, D. (2012) Diversity and technological progress. In *The Rate and Direction of Inventive Activity Revisited*, pp. 319–360. University of Chicago Press.

Acemoglu, D. and Autor, D. (2011) Skills, tasks and technologies: implications for employment and earnings, *Handbook of Labor Economics*, **4**: 1043–1171.

Acemoglu, D. and Robinson, J. (2012) *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. New York, NY: Crown Busi.

Acemoglu, D. and Restrepo, P. (2018a) The race between man and machine: implications of technology for growth, factor shares and employment, *American Economic Review*, **108**: 1488–1542.

Acemoglu, D. and Restrepo, P. (2018b) Robots and jobs: evidence from US labor markets, *Journal of Political Economy*, doi:10.3386/w23285

Acemoglu, D. and Restrepo, P. (2018c) Modeling automation, *AEA Papers & Proceedings*, **108**: 48–53.

Acemoglu, D. and Restrepo, P. (2018d) Artificial intelligence, automation and work. In *The Economics of Artificial Intelligence: An Agenda*, pp. 197–236. National Bureau of Economic Research, Inc.

Acemoglu, D. and Restrepo, P. (2019) Automation and new tasks: how technology changes labor demand, *Journal of Economic Perspectives*, **33**: 3–30.

Agarwal, A., Gans J. S. and Goldfarb, J. S. (2018) *Prediction Machines: The Simple Economics of Artificial Intelligence*. Cambridge MA: Harvard Business Review.

Allport, G. W. (1937) *Personality: A Psychological Interpretation*. Holt and Co.

Autor, D. H. (2015) Why are there still so many jobs? The history and future of workplace automation, *Journal of Economic Perspectives*, **29**: 3–30.

Autor, D. H., Levy, S. M. and Murnane, R. J. (2003) The skill content of recent technological change: an empirical exploration, *The Quarterly Journal of Economics*, **118**: 1279–1333.

Ayres, R. U. and Miller, S. M. (1983) *Robots: Applications and Social Implications*. Ballinger Publishing Co.

Bessen, J. E., Impink S. M., Reichensperger, L. and Seamans, S. M. (2018) *The Business of AI Startups*. BU School of Law, Law & Economics Series Paper No. 18-28. Available at: https://ssrn.com/abstract=32932.

Cassidy, S. (2004) Learning styles: an overview of theories, models, and measures, *Educational Psychology*, **24**: 419–444.

Dauth, W., Findeisen, S., Südekum, J. and Woessner, N. (2019) *German Robots—The Impact of Industrial Robots on Workers*. Institute for Employment Research Discussion Paper.

Dosi, G. (1982) Technology paradigms and technological trajectories, *Research Policy*, **11**: 147–162.

Dreyfus, H. L. (1992) *What Computers Still Can't Do: A Critique of Artificial Reason*. MIT Press.

Forester, T. (1985) *The Information Technology Revolution*. Cambridge, MA: MIT Press.

Graetz, G. and Michaels, G. (2018) Robots at work, *The Review of Economics and Statistics*, **100**: 753–768.

Groover, M. P., Weiss, M., Nagel, R. N. and Odrey, N. G. (1986) *Industrial Robotics: Technology, Programming, and Applications*. McGraw-Hill.

Harari, Y. N. (2018) *21 Lessons for the 21st Century*. Spiegel & Grau.

Honey, P. and Mumford, A. (1986) *The Manual of Learning Styles*. Peter Honey Associates.

Judis, J. B. (2016) *The Populist Explosion: How the Great Recession Transformed American and European Politics*. New York, NY: Columbia Global Reports.

Kellner, T. (2018) Game On: Augmented Reality Is Helping Factory Workers Become More Productive. Available at: https://www.ge.com/reports/game-augmented-reality-helping-factory-workers-become-productive/ [Accessed 19 April 2018].

Lanier, J. (2018) *Ten Arguments for Deleting Your Social Media Accounts Right Now*. New York, NY: Hoffmann and C.

Mantoux, P. (1927) *The Industrial Revolution in the Eighteenth Century: An Outline of the Beginnings of the Modern Factory System in England*. London, UK: Harcourt.

Manuel, R. and Castańeda, A. (1974) *Cultural Democracy, Bicognitive Development, and Education*. Academic Press.

Mazzucato, M. (2015) *The Entrepreneurial State*. Anthem Press.

Mokyr, J. (1992) *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford University Press.

Muralidharan, K, Singh, A. and Ganimian, A. J. (2019) Disrupting education? Experimental evidence on technology-aided instruction in India, *American Economic Review*, **109**: 1426–1460.

Neapolitan, R. E. and Jiang, X. (2018) *Artificial Intelligence: With an Introduction to Machine Learning*, 2nd edn. Chapman and Hall/CRC.

Nelson, R. and Sidney, W. (1977) In search of useful theory of innovation, *Research Policy*, **6**: 36–76.

Nilsson, N. J. (2009) *The Quest for Artificial Intelligence: A History of Ideas and Achievements*. New York, NY: Cambridge University Press.

Olmstead, A. and Rhode P. (2001) Reshaping the landscape: the impact and diffusion of the tractor in American agriculture, 1910–1960, *The Journal of Economic History*, **61**: 663–698.

Ong, S. K. and Nee A. Y. C. (2013) *Virtual and Augmented Reality Applications in Manufacturing*. New York, NY: Springer Science and Business Media.

Pasquale, F. (2015) *The Black Box Society: The Secret Algorithms that Control Money and Information*. Cambridge, MA: Harvard University Press.

Rasmussen, W. (1982) The mechanization of agriculture, *Scientific American*, **247**: 76–89.

Russell, S. J. and Norvig, P. (2009) *Artificial Intelligence: A Modern Approach*, 3rd edn. Prentice-Hall.

Stiglitz, J. (2012) *The Price of Inequality: How Today's Divided Society Endangers Our Future*. W. W. Norton & Company.

Zeira, J. (1998) Workers, machines, and economic growth, *Quarterly Journal of Economics*, **113**: 1091–1117.

Zuboff, S. (2019) *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. Public Affairs.