

THE CONSEQUENCES OF AI-BASED TECHNOLOGIES FOR JOBS

Peter Cappelli

Professor of Management, Director of the Center for Human Resources at the Wharton School of the University of Pennsylvania and Research Associate at the NBER

Summary

This contribution follows the recent public debate on the changes across industrial countries that stem from information technology, including notions of artificial intelligence and its implications for how work is performed. While acknowledging the size and pervasiveness of these discussions, the article discusses the core arguments related to the impact of information technology on the way businesses and organisations operate, how these changes could translate to the labour market, and other potential outcomes such as lower wages or unemployment.

The argument begins with an introduction to the two ways in which people tend to anticipate future developments. This either happens through estimates based on prior experience (commonly known as forecasting) or through a belief in a real uncertainty of future developments and reliance on other kinds of evidence besides

traditional forecasts. The article maps the projected impact of technological uptake on the labour markets and reviews the empirical evidence. It touches upon many of the above-discussed trends, such as skill-biased technological change or routine-biased technological change, and their implications for skills demand. Applying an historic perspective, the article argues that predictions based on the past may be less relevant in the current context. Although new equipment and practices could eliminate certain jobs, on balance they do not necessarily destroy jobs because their overall effects on improving productivity and overall wealth create jobs elsewhere.

To understand why assumptions claiming that the future is like the past are not correct and extrapolations from prior experiences are unlikely to be accurate predictors of the future, read this chapter.

Recent discussion first in the business press and then in related public policy communities has considered the notion that industrial countries are on the verge of important changes that stem from information technology (IT) broadly, including notions of artificial intelligence (AI), and its implications for how work is performed. The size and pervasiveness of these discussions

merits a serious look at the ideas behind them and the fundamental question they ask: is there something happening already or about to happen in information technology that will change in a fundamental way businesses and organisations, jobs, and outcomes like pay and unemployment? I consider these issues below.

1. The nature of the discussion

Before considering the arguments and assertions about the implications of evolving IT, it is worth thinking through the context in which those stories take place. Followers of the

media are well aware that there is a bias toward reporting stories that represent something new, especially something new and dramatic. That includes claims about developments that will

happen, even if there is little or no evidence of them yet. We may notice these stories especially when they relate to health, e.g. epidemiological studies showing that some particular food group is associated with either remarkably better or worse life outcomes. It is extremely difficult to run a story that says, for example, 'still nothing new in effective weight loss'. A first question to ask is whether the apparent magnitude of the stories of technological change reflects a change in the nature of the media and public discourse rather than reflecting something about the merits of the arguments themselves.

There have been changes in the media that might help create the impression that particular stories are more important than would have been the case in the past, such as the fact that there are now many more outlets for stories, including

social media, where surprising or frightening accounts are repeated and reinforced over and over. There is also considerable expansion of organisations focused on public policy, especially those businesses which advocate ideas that are important and support those that attract attention. Hosting discussions, producing reports, commenting on media stories are standard practices for such organisations. Every major consulting company now produces reports and markets their views on policy-related stories, including technology and workplace topics.

The fact that there is a great deal of discussion about IT certainly suggests that it is a topic worth investigating, although it is not *prima facie* evidence that the arguments which provoke that discussion are correct. The truth is typically more boring than the speculations.

2. Anticipating the future

Assessing the merits of arguments about the potential effects of IT in the workplace or elsewhere should begin with thoughts about epistemology: what is it that we know, and how can we know it? Specifically, how can we distinguish reasonable belief from mere opinion? What constitutes knowledge is always a pertinent question, but it is especially important in this context because of the unique nature of the claims being made. They are claims about the future rather than the present, although they may well be informed by the present.

There are at least two quite different types of claims about the future that are made in the social sciences. The first concerns probabilities and risk: we have very little idea about, for example, whether my house will burn down but, based on prior experience of houses like mine, we can estimate with considerable accuracy what the odds of that are.

Forecasts move us from predictions about common events and about individual units in a population to anticipating events that have not happened before. They go a step further than identifying average experiences in the past to extrapolate from the past. To predict, for instance, the unemployment rate in a year's time, they look back to previous unemployment rates and to variables that determined them or at least were associated with them. If the model using those variables explained a reasonable amount of the variation in previous unemployment rates then we will try to use it to extrapolate into the future. We do so by assuming that the structure of the model remains the same going forward or, in practical terms, that the coefficients of regression-related models in the future will be the same as they are in the model. Assuming we have more recent values for the variables in the model, we apply them to that model and

generate an estimate or forecast as to what the unemployment rate will be in the future.

A great advantage of this approach in terms of epistemology is that we have some ability to assess how accurate our forecast of the future is, based on how well our model has predicted outcomes in the past.

The downside of the approach is that the assessment of accuracy does not work, nor will the model produce an accurate forecast, if the model's underlying structure (the relationship between the variables and the outcome being forecast) changes from the earlier period. For example, economic forecasting models in the United States that proved remarkably predictive in the 1960s stopped being very accurate in the 1970s and after, apparently because of changes in the structure of the economy. It took some time to recognise that change, and the accuracy of the models never recovered to their previous levels.

The second type of claim is one where we believe that there is true uncertainty about the future, where average experience in the past is not likely to continue into the future, and the structure of forecasting models changes in ways that are not clear *a priori*. In this context, the concerns of epistemology become much more important. Other kinds of evidence besides traditional forecasts also become more important. For example, explanations that have predicted well in the past, perhaps in different contexts, might be useful. The role of theory that has been supported over time by evidence becomes important. We might not now have a good idea what the effects of new technologies will be in the future, for example, but we might well believe that the effects of previous technologies would be informative and that principles like supply and demand will still be relevant in explaining what they will be. Other evidence might include examples consistent with the prediction in subsets of the population

or trends in the direction of the prediction (e.g. leading companies are doing this).

The complication in assessing claims about the influence of IT and AI is that most of the attention-getting claims are based on the assertion that the future is not like the past, that the new developments in AI will change the structure of the relationships such that extrapolations from prior experiences are unlikely to be accurate predictors of the future. We might think of this as a double uncertainty: we cannot say with any certainty what IT innovations will look like in the future, let alone how they will affect the economy. Such claims are difficult to assess in traditional ways because they do not have an empirical basis. When we cannot test how well explanations actually work – in this case because the events being explained have not yet happened – we are forced to use other kinds of assessments.

These other approaches rely on the structure of the arguments being made. A common standard is whether the explanations are deduced appropriately from principles that have already been established, the standard deductive-normative format for generating normal science hypotheses. Beyond that, we often use criteria that are not well justified, such as *ad hominem* arguments – the person making the case has been right before or they are an 'expert' on the topic.

In recent decades, one of the more important developments in business has been to come to grips with the problem of uncertainty. On the one hand, we can never be certain about any aspect of the future, although we may be confident that some aspects are good enough to plan on, such as the sun coming up tomorrow. But what can we do when we are aware that our predictions or forecasts are not very good? We use these concepts below to consider the merits of the arguments about the future impact of AI.

3. The nature of the claims

A major complication in assessing the claims about what AI might do to the workplace and to employment in particular is that there are so many of these claims. In many cases, the same individuals have made quite different claims over time, requiring some condensing and organising of them.

The place to begin is with a definition of AI. A standard for determining what a term means is that it should not overlap with other terms that refer to similar concepts. The common and arguably standard definition of AI in dictionaries and elsewhere dates from a 1956 symposium of cognitive scientists who proposed a research programme to investigate it (Minsky, 1994). The general idea at the time was that AI is machine-based thinking that mimics what humans can do.

However, defining what thinking actually is continues to be elusive. Alan Turing (1950) proposed a simple test of AI which is whether a machine could fool a person into believing that its responses to questioning were actually from a person. More sophisticated and didactic definitions focus on thinking that only humans can do, which includes reasoning, judgment and learning. By that definition, AI would seem to be a continually shrinking domain as machines become capable of more and more tasks: computing power and programs that formalise decision-making enable computers to solve more problems. Calculations that only humans could do generations ago can now be done on pocket calculators.

Definitions of AI continue to change as practice changes. At least some observers have abandoned the notion that AI is about distinctly human intelligence and describe it as the study of any kind of intelligence; others differentiate between 'weak' or 'narrow' AI, focused on solving

particular problems, and 'strong' or 'general' AI that can solve problems across domains.

Whether one sees these debates over the nature of AI as semantic, practical, reacting to developments in practice, or conceptual – ultimately turning on epistemology and notions of knowledge – securing agreement on a definition is difficult. Fortunately, it is probably not necessary for the task at hand to have a clear differentiation about what AI means as the claims about effects on the labour market are mainly about IT as it is conceptualised now.

Arguably, the most useful applications of computers today are in data science with the most immediate implications for jobs. Here, many of the new applications do not necessarily involve reasoning, judgment, learning or anything like thinking. 'Big data', for example, is simply software to handle statistical processes with data sets that had been too large for traditional programs to handle; machine learning, at least in its general format, is a technique for finding relationships between variables; and algorithms are just decision rules derived from evidence that do not necessarily require computer power, while those derived from machine learning make predictions that can be validated. Natural language processing and speech recognition are, in essence, pattern-recognition problems that become possible for machines to do as computing power increases. Most of the claims concerning the effects of AI are, in fact, assertions about data-science tools like those above.

The next step in beginning our analysis is much more straightforward: i.e. to consider the outcomes of IT that are of interest. Following the debate in the popular press, we are concerned with the effects of AI on jobs – in particular, whether it increases or reduces the

number of them – and, to a lesser extent, how it might change the tasks required of jobs, the skills needed to perform them, and the quality of jobs widely considered.

That leads directly to the claims about the effects of AI that are currently the focus of atten-

tion. The most important of these are assertions that developments in AI will eliminate large numbers of jobs and, in the process, create long-term structural unemployment and lower wages, especially for lower-skilled individuals.

4. A brief history of research on AI and the labour market

Concerns that modern technology will lead to unemployment go back to the early days of industrialisation, at least to the Luddites in the early 1800s who protested against the new factory system that threatened the income of more skilled workers (Thomas, 1970). The possibility that new industrial technology was eliminating jobs became a long-standing political question thereafter in the UK and in much of Europe, but less so in the United States where unemployment, at least until the Great Depression, was less of a concern. In all industrialising countries, the mechanisation of farming, along with new agricultural techniques, were displacing workers and the concern arising from looking at projections was that the manufacturing economy could not accommodate all those soon-to-be displaced workers (Fano, 1991). The Great Depression kindled the debate about the role of technology in jobs, not just because unemployment was so high but because the evidence even then suggested that, in the 1930s, the United States experienced a massive jump in productivity (Bix, 2000) that was seen as contributing to job losses.

Nevertheless, in the 1960s, a period of dramatic economic growth and low unemployment, concern that technology and automation were causing unemployment was a political concern because of the perception that technology was and would be advancing quickly. America's

President Johnson set up a commission to investigate the evidence for that concern, which subsequently concluded that there was little evidence for it (Automation Commission 1966).

The concern about computers and jobs *per se* developed later, partly because the rise of computers became quite gradually. Perhaps ironically, initial concerns appeared to be driven by a question of financial accountability when investments in computers and IT generally began to increase. Complaints from the world of investors questioned these investments because there did not appear to be an associated pay-off from them in terms of operating efficiencies (e.g. Straussman, 1997). The famous quip from economist Robert Solow – ‘we can see the computer age everywhere but in the productivity statistics’ – captured the difference between the rhetoric about the value of IT and the apparent reality. That apparent reality became known as ‘the productivity paradox’.

For our purposes, the evidence on IT investments and productivity matters because productivity is typically measured in terms of labour, output per employee. The most straightforward manner in which productivity increases is when firms use fewer workers for the same output, or a smaller proportion of workers for greater output. Dedrick et al. (2003) review the earlier literature on

this topic and note that initial studies, through the mid-1990s, did not find evidence of any significant return on the investment in IT.

Research into labour economics about computers had been energised by Krueger's (1993) finding that wages were higher, other things being equal, for individuals who used computers at work. This finding helped to kick off a number of arguments that are continuing today, suggesting that using computers contributes to better-paying jobs, presumably because such jobs require more skill. (It should be noted that this is the opposite of most contemporary claims that computers will make outcomes worse for workers.) The implications were that jobs that did not require computers would fall behind in pay, helping to explain an aspect of the 'digital divide', inequality of various kinds but especially in pay associated with access to IT and the internet.

Cold water was thrown on this conclusion – although frankly only in the academic world – by DiNardo and Pischke's (1997) finding that workers who used pencils also earned higher pay. Their tongue-in-cheek title about pencil use referred to their finding that workers who were using tools associated with working at a desk earned more, suggesting that it may not have been the use of computers that was associated with higher wages but simply doing the kind of jobs for which computers would be useful that paid off. The study illustrated the common problem of omitted variables, in this case that what was associated with computer use also mattered.

By the 1990s, there were two different streams of research interested in the relationship between computer use and employment outcomes: economists studying the effects of IT on business, whose interest was looking for productivity improvements, and economists and some sociologists, whose interest was looking for explanations for wage differences.

As Dedrick et al. (2003) note, the former stream of research shifted for the analysis from the national and industry-level down to individual firms where they began to find evidence of greater business outcomes associated with IT investments. These results were replicated in Europe although not in developing countries, while the size of the effects appeared crucially to depend on accounting decisions that determine which costs are associated with IT investment measures: is it just the hardware and software, does it include the training costs of employees, the reorganisation costs, and so forth.

An important finding in many of these studies was the considerable variation in the relationship between IT and performance across organisations. Bresnahan (1999) helped kick off a new direction in the IT productivity debate related to that variability by focusing on the changes in business organisation – more commonly referred to today as restructuring – that are associated with the successful introduction of IT investments. Bresnahan, Brynjolfsson and Hitt (2002) and a string of subsequent studies identified the synergies between investing in IT and changing the organisation of work to explain performance improvements. This research relates to the DiNardo and Pischke notion that it may not be the computers *per se* that are driving the outcomes of interest but rather the changes in existing practices that they produced.

On the labour economics front, Autor, Katz and Krueger (1998) found that skill upgrading was greatest in those industries that had made largest investments in IT, suggesting a different complementarity between labour and IT. This result is related to the earlier Krueger (1993) finding – the idea that computer use raises skill requirements and, in turn, wages. Autor, Levy and Murnane (2003) examined the apparent association between the introduction of computer-based systems and more college-based labour with an explanation that computers take over repetitive, lower-level tasks and

therefore eliminate lower-paid jobs and provided evidence at the economy level to support it. These studies align with others about the rising relative wages of college graduates compared to those with qualifications less than college degrees to reinforce a notion that became known as ‘skill-biased technology change’. This view of the world articulated by the labour economic studies remains dominant in the popular press although, as is shown below, the evidence related to it increasingly counters that view.

Before turning to the extensive body of research carried out since then, virtually all done in economics, it is important to understand some of the assumptions that underlie that research. First, when economists talk about ‘technology’ in the broad sense, they mean anything that changes the production function – new management techniques, capital investments in equipment or IT, presumably even new priorities, and so forth (see Auto, Katz and Kearney 2008 for an explicit statement on this). Observers often assume that conclusions about the effects of technology refer to IT, but unless the studies are measuring IT explicitly, that is not the case.

Second, with few exceptions, studies that measure computer investments claim to be capturing the influence of IT *per se* and not, as Bresnahan and others found, a mix of organisational transformation and new ways of organising work which are associated with the introduction of computers. This relates to the *ceteris paribus* assumption and, when it is violated, to the problem of omitted variables.

Third, the assumption is that the educational qualifications of those in jobs are an accurate measure of the requirements of those jobs. The practical reason for this assumption is that it is relatively easy to access data on the education of individual employees but quite difficult to get data on the requirements of jobs. As a result, changes in the percentage of individuals with college degrees and in the wages associated with those jobs are interpreted as changes in skill requirements and in the demand for skill. Careful observers, especially those outside economics, question the reasonableness of that assumption (see, e.g. Liu and Grusky, 2013).

Finally, economists, indeed all social scientists, attempt to advance arguments associated with their paradigm typically at the expense of other explanations. It is often heard that historians attempt to provide a complete explanation of the phenomena they are studying, but there is no credible claim for that in the social sciences. A simple explanation, consistent with the underlying paradigms, is far preferable in our respective disciplines to a complicated explanation that includes multiple and particularly unrelated components, even if the latter explains much more of the phenomenon. Evidence for this is easy to see in any empirical study, where the amount of variation explained by the explanations submitted is only a fraction of the total variation.

This last point is especially important in making sense of the research on IT where it is often claimed that x is the cause of y when, in fact, the best we can claim is likely to be that x is one factor associated with y.

5. Skill-biased technological change

Although not related to IT *per se*, the notion of skill-biased technological change is often used to explain or at least support the claims about

how IT is changing outcomes in the job market. At its heart is an older theoretical argument often credited to Polanyi (1944) which asserts

that new technology inevitably raises skill requirements, because higher skills are needed to use the new technology¹. The inevitability assertion is manifestly not true as the thrust of modern industry and techniques such as scientific management were designed precisely to reduce the skill requirements in individual jobs, e.g. by breaking them up into simpler sub-tasks. (It may well be true that the initial introduction of a new technology, such as computers, requires considerable skill to use them, but later modifications make them easier and easier to use. For example, cash registers with pictures on them are computers for checkout assistants that do not even require literacy. 'Technology' in these studies is not measured directly but is assumed as an underlying development of modern economies.

Katz and Murphy's (1992) extremely influential study arguably kicked off the contemporary version of this idea by finding that the 'college premium' – the ratio of what an average college graduate earned in the economy to what the average high-school graduate earned – rose sharply in the United States at a time when the proportion of the labour force with a college degree was also rising. Despite the rising supply, the apparent price of skill had also been rising, as measured by the college wage premium. The authors argued that changes in demographics and, more generally, on the supply side did not account at least for the recent rise in the college premium, so the explanation must lie with an increase in demand.

Something of a consensus developed among many that new technology, particularly information technology, caused an increase in the demand for skill. The topic was particularly popular because it was seen as an explanation for the dominant issue of the early 2000s, which was rising wage inequality. Many studies followed the Katz and Murphy paper in exploring

changes in the college wage premium. A broader and more general study of the relationship between education, technology and wages makes a similar claim over a much longer period of time, suggesting that surges in the supply of college graduates moderated the fairly continuous increases in the demand for skill in American economic history (Goldin and Katz, 2008). The phrase 'skill-biased technological change' emerged from these empirical studies.

Although they received less attention, many studies questioned the skill-biased technological change idea. In particular, the occupational shifts that seemed to be the basis of the evidence of skill upgrading had been under way for at least a decade before IT investments became substantial. Card and DiNardo (2001) noted that the college wage premium did not track measures of actual technological change well and concluded that it was not a very helpful concept for understanding changes in wage structures. Card and Leimuix (2001) found that, in the 1990s, the sharply rising college premium was not true across the labour force but was mainly attributable just to the experience of young people. (Mishel and Bernstein (1994) present a sweeping critique of the IT explanation.)

Despite the lack of correspondence with much of the evidence, skill-biased technological change had a great deal of appeal because it was useful in understanding growing wage inequality, a topic of enormous policy interest, and the related issue of the apparent growing wage premium for college graduates over non-graduates. Later critiques further weakened empirical support for the idea, however. Schmitt, Shierholz and Mishel (2013) presented a series of examples in which the notion of skill-biased technological change is inconsistent with the evidence. This included the fact that it was

1 Polanyi actually says very little about technology as his arguments focus on the relationship between markets and institutions in the transition to industrial economies.

inconsistent with wage trends after 2000. More recently, Beaudry, Green and Sand (2014) found that the demand for higher skill appears to have declined since the early 2000s. Valetta (2017) also found that the college premium has been declining.

Acemoglu and Autor (2012) signalled a pivot away from the simple view of skill-biased technological change. They noted that it did not work outside of the 1963-1987 period which was the basis for the Katz and Murphy study. They calculated workers' average weekly, with inflation discounted, over time and by education level – high-school dropouts, high-school grads, those with some college background, college grads, and those with graduate degrees – and found that the wage gap between those different groups in the early 1960s and then again in the mid- to late 1970s was quite small, as Richard Freeman had noted earlier. Then right after the 1981 recession, real wages for everyone with less education than a four-year college degree started to collapse and continued to decline through the early 1990s. The rapid decline in high-paying, union manufacturing jobs and the rise of low-wage competition from China in particular certainly played a big part in the explanation. Although wages for college grads did not take off, they did eventually recover some of their lost ground.

The result of these two movements – the decline of real wages for everyone, the continuing decline for high-school graduates, and the modest improvement for college graduates – created the wide gap between the groups and a sizeable wage premium for college graduates which started in the 1980s. The fact that the college premium appeared to be caused more by the decline in high-school wages than by the rise of college wages did not fit the demand-side explanation of skill- biased technological

change. It appeared to be a story about which group lost the fastest as both high-school and college graduates have seen a fall in real wages since 2002 (Shierholz, Davis and Kimball, 2014). Demographic trends also had a big effect on wages across age cohorts (Jeong, Kim and Manovskii, 2014) which affected the college premium across cohorts; the college premium for students from poorer families is about half of that for wealthier families (Bartik and Hershbein, 2018), partly reflecting the graduates' unobserved attributes. More than one-fifth of the college wage premium also appears to be associated with cost of living differences because college graduates tend to live in more expensive places than high-school only graduates (Moretti, 2011).

Acemoglu and Autor (2012) moved the discussion back towards a different explanation of technology that was consistent with Autor's earlier studies – i.e. that computers in particular eliminate routine jobs. The difference now is the assertion that those routine jobs were in the middle of occupational and wage structures. We could call this the 'hollowing out' view. From this point on, most research abandoned the simple notion of skill-biased technological change that economic growth inevitably generated higher skill requirements.

Schmitt, Shierholz, and Michel (2013) presented a sweeping critique of the hollowing-out notion as well, noting that it does not explain changes in occupational distribution after 2000 (in particular, low-wage jobs have been growing), that occupational changes have not driven changes in the wage distribution, and perhaps more importantly, that changes in the occupational distribution associated with a shrinking middle began long before the modern computer age². Barany and Siegel

2 The fact that the studies from the Economic Policy Institute (EPI) attack so consistently the simple explanations for changes in wages and jobs may be seen by some as reflecting an interest in focusing the discussion on the role of policy in shaping labour market outcomes. However, but it is also fair to note that, unlike the paradigm-based research articles, they are focused on explaining the phenomena per se rather than advocating a conceptual explanation.

(2018) document that the declining middle in the US occupational structure was under way decades before the IT expansion of the 1990s and appears to be related in the economy as a whole to the shift from manufacturing jobs to service jobs. We consider more studies below on IT *per se* that also contradict this notion.

What should we conclude about the skills-biased technological change idea? First, the original incarnation of the argument, that technology inexorably increases skill requirements and, in turn, alters the demand for skill and wages, has been largely abandoned by researchers. Second, the job-polarisation version differs fundamentally from the original – in particular, there is no assumption of ever-increasing skill requirements – and mainly only shares an underlying supply-and-demand mechanism.

As Howell and Kalleberg (forthcoming) note in their extensive review of explanations for recent wage and occupation changes, there are other

explanations at least equally – and arguably more – compelling than job polarisation for labour market outcomes. These focus on changing power relationships which have allowed employers to squeeze lower-skilled workers and the highest earning individuals to secure more income. For example, Kristal (2013) finds that the introduction of computers made workers more replaceable which lowered their wages. These arguments do not have the advocacy the job-polarisation idea and its supply-and-demand underpinning have, at least among a large number of economists studying labour market outcomes.

As shown below, there is certainly some evidence for IT changing occupational structures, although how much of the change is truly driven by IT as opposed to coinciding with trends already under way, and how much is caused by factors associated with IT, such as the associated restructuring of organisations, is not clear.

6. Forecasting the effects of IT on jobs

Although the above-mentioned research has had considerable influence on popular thinking about the effects of IT, more important for our purposes are the studies concentrating on the topic of IT use. Recently, much and arguably most of the research on the relationship between IT and jobs has been motivated by the practical concern as to whether IT will eliminate jobs. This stream of research has been motivated largely as a reaction to forecasts, specifically pessimistic forecasts, about the likely effects of continuing advances in IT which claim that new and emerging developments in computing power, in software, and in data science are fundamentally different from those seen before.

Arguably the most important of these prediction arguments is from Brynjolfsson and McAfee (2011) who argue that the IT technology emerging now is fundamentally different from what has been seen before and will affect the workplace differently than what has been seen before. The most attention-grabbing claim in their book, which appeared at a time of substantial unemployment in the United States, is that this new technology will lead to substantial job loss. Schwab (2016) essentially adopted this view, as did many reports written by consulting companies.

It is not possible here to review or even catalogue all the reports from outside the academic and policy world, although they

have some common themes. First, in terms of approach, they are typically authored by practitioners outside IT fields. They tend to rely on surveys that ask executives what they believe about the future. Second, in terms of conclusions, none of them appear to claim that the future will look more or less like the past or that the changes associated with IT are similar to those experienced before. The typical conclusions repeat assertions that IT will 'disrupt' the way business is done and that businesses need to figure out how to deal with these developments. Many of these conclusions are dramatic: Bain, for example, forecasts that half of all current jobs in the United States could be eliminated in 15 years and that US employers will need 30 to 40 million fewer workers by 2030 (Harris, Kimson and Schwedel, 2018).

By contemporary research standards, these claims contradict evidence which has been consistent since the Industrial Revolution that while new equipment and practices eliminate certain jobs, on balance, they do not destroy jobs because of their overall effects on improving productivity and overall wealth create jobs elsewhere. Autor (2018) articulates the many paths through which technology that increases productivity boosts economic growth and why, in modern history, it has not yet led to job losses.

As noted above, the epistemological problem raised in assessing these reports is how to separate assertions that we might dismiss as mere opinion from something that we would consider a true belief. If it is reasonable to conclude that future developments in IT are so unlike the past that we cannot use prior experience to assess them, then we cannot use evidence to assess those assertions.

One approach, adopted below, is to dismiss the claim that when new IT developments come they will be so distinctive that we cannot learn anything about their likely effects from prior

experience with technology. When we think about historical developments in transformative technology, such as the rise of steam power, electricity, the first computers, and so forth, it does not seem credible to suggest that nothing could be learned from such experiences. If we have yet to see these technologies, then assertions about whether their effects will be so different from anything seen before seems very much like opinion rather than a true belief.

There are areas of inquiry where predictions are made consistently about events for which we cannot generate traditional forecasting models because in the past there were not enough similar circumstances – possibly none – to use as a basis. We could consider these *sui generis* predictions. Concerns about how a political leader will react to a challenge, whether countries will go to war at a particular moment, or whether 'society has changed' may fit this prediction category. It is also the case that we have to make predictions where forecasting models are at least conceptually possible although, for a variety of reasons, such as time pressure or lack of resources, they cannot be constructed.

We might describe the effort to make such predictions as 'expert judgment'. Tetlock (2017) studied the phenomenon of predictions by experts extensively, in particular with respect to political events. He found that experts' accuracy in making these predictions barely surpassed 'monkeys tossing darts at a dartboard' or, less creatively, were no better than chance. Predictions of societal and political events are perhaps not common enough to be able to tell if those who are 'good' at predicting have just been lucky. However, Tetlock and Gardner (2018) engaged in a sizeable exercise to see what makes some individuals better than others at actually predicting events that could be confirmed later. Their conclusions are important to bear in mind when looking at forecasts concerning the future of IT.

Those who are worse at predicting are highly confident of their abilities – over-confident; experts who are deeply focused on their subject, ‘hedgehogs’ according to Isaiah Berlin, are also worse when compared to those with wider expertise, the ‘foxes’. Followers of grand theory, which would include the economics paradigm, are worse at predicting. Conversely, those who question assumptions, who look for comparable situations and events elsewhere, and who consider the counter arguments to their positions do better at predicting.

The reports above tend to assume the most important conclusion – that IT developments will be transformational – and from there pursue implications that sometimes extrapolate from current circumstances. Applying Tetlock and Gardner’s (2018) criteria, the studies rarely, if ever, question or even identify their assumptions, consider counter arguments, or believe that much could be learned from other circumstances. It is also worth noting that consulting companies in particular have a material interest in securing business that is not always perfectly aligned with presenting the most accurate story. These reports are marketed aggressively and have considerable influence on business leaders who, in turn, are often the empirical source for the next set of studies.

One of the most influential predictions about the impact of IT, especially among practitioners, was conducted by Frey and Osborne (2017). It asked computer experts to assess whether, under the best circumstances, it was possible for computers to take over the central tasks of a set of jobs or if it will be possible to do so soon. Their assertion that almost half of the jobs could be taken over by computers forms the basis for the conclusion in many of the practitioner reports that those jobs *will* be taken over by computers and soon.

Unfortunately, the prediction stopped there. The question did not ask for a prediction of what will actually happen in the real world. There is

an enormous gap between what is technically possible to do, the question asked of computer experts, and what is practically useful or financially viable to do. We can, for example, go to construction sites almost anywhere in the world and find tasks being performed by hand that could easily be performed by existing machines. The fact that loads are carried by hand and holes dug using shovels in many parts of the developing world reflects the fact that labour is so much cheaper than equipment, not that the workers are unaware of trucks or backhoes. Then there are tasks that IT and robots can perform now, although they are not good at them. Mechanical robots can create alcoholic mixed drinks the same way as bartenders do, but a colleague who observed this indicated that the quality of the drinks was poor and it took two employees to support and service the robot whenever it was in operation. The machine did the task, poorly, and at incredible expense.

The Organisation for Economic Co-operation and Development (OECD) (2018) took the Frey and Osborne estimates at face value and then used estimates of job requirements from the Programme for the International Assessment of Adult Competencies (PIAAC) skills survey and concluded that roughly 14% of jobs met the criterion that machines could or soon would be able to perform them – i.e. A much smaller number. Whether they will take over those tasks and whether doing so will eliminate jobs is another question considered below. Arntz et al. (2016) had earlier conducted an estimate similar to that of Frey and Osborne and concluded that 9% of employees were in jobs that were likely to be automated.

Forecasts for the effects of technology have been more difficult to predict than the political and social events studied by Tetlock (2017) and Tetlock and Gardner (2018). In fact, there is something of a sport in reminding us of how poorly we have been able to anticipate not only which technologies will succeed and when they

will arrive but what their influence will be when they do. For example, Funk (2017) revisited the technology predictions of MIT's Technology Review and found few examples of success, while management scholar Joseph Switter (1965) predicted that, by 1985, computers would take over most management tasks. Predicting the implications of technology was a hot topic in the 1960s, when researchers were aware of the many factors outside of

technology *per se* that affect its introduction, such as actual demand for it, especially relative to competing solutions, social and political implications of using the technology, and so forth. They articulated techniques for making such predictions that include analysing switching costs to new technologies (see Quinn 1967 for an example), none of which seem to be used in the current forecasts.

7. Evidence of the effect of IT on jobs

We turn now to recent empirical evidence that relates to the predictions above. Beginning with the Brynjolfsson and McAfee (2011) assertion that new IT technology is fundamentally different and will lead to a net reduction in jobs, the current economic environment, at least in the United States with record low unemployment, offers that notion less sympathy than when it was articulated during the Great Recession. More recent research gives it no support. The job-polarisation hypothesis – that IT is eliminating and will continue to eliminate more routine jobs – also receives little support in more recent research.

Bessen (2016) looks at US data and finds that increased IT use is actually associated with more jobs. He also finds no evidence of job polarisation associated with greater IT use. Aum, Lee and Shin (2017) found that IT investments were actually smaller for lower-level jobs doing routine work than for higher-level jobs, which is inconsistent both with an earlier view that IT eliminates lower-level jobs and with the notion that it disproportionately targets middle-level jobs. Gregory, Salomons and Zierhn (2016) also conclude that there is no evidence of IT use reducing employment in Europe. Boreland and Coelli (2017) examine IT use and employment in Australia and find

no evidence that greater IT use has reduced employment or has it decreased employment in jobs that would seem to be routine in terms of skill. In fact, there is no evidence that greater IT use has been associated with greater changes in sectors of the economy where IT investments have been the greatest.

The underlying logic behind the job-loss idea is that where IT does not eliminate jobs altogether, it changes skill requirements, rendering incumbents unqualified for further employment and costing them their jobs. Allen and de Grip (2012) examine the general question of whether skill obsolescence increases the probability that individuals will lose their jobs and conclude that, in practice, it does not. One explanation for that lack of relationship is that individuals and employers recognise when skills may become obsolete and respond accordingly, through retraining and other ways.

An important issue in understanding the outcomes of IT on jobs and labour outcomes in general is the distinction between tasks and jobs. Jobs are typically defined as a collection of tasks. Except for the very simplest assembly-line work, most jobs include many tasks: virtually every job description and employee handbook in

the United States ends the description of any job with the phrase ‘and tasks as assigned’, which means that supervisors can add virtually any task to the job of any employee.

This simple fact that jobs comprise many tasks gets to the heart of many misunderstandings about the effects of IT on employment. The applications of IT to work are typically task-by-task: at the lower-skill end, dispensing cash through ATMs, at the higher end, reading x-rays and digital images. The reason the pundits were wrong in expecting that ATMs would eliminate bank teller jobs is that tellers have many tasks besides simply dispensing cash. Radiologists do read x-rays, but they also have many other tasks, including consulting with other doctors and patients, advising on treatment, and so forth, which means that algorithms which ‘read’ x-rays do not eliminate their job (Brynjolfsson, Mitchell and Rock, 2018 acknowledged this complication).

The fascination with autonomous or self-driving vehicles that swept the business press a few years ago fixated on the prediction that such vehicles would eliminate the job of truck driver: the European Automobile Manufacturers’ Association (2017) predicted, for example, that half of all trucking jobs would be gone within 10 years. That conclusion ignored the reality of what most truck drivers do, which is to make deliveries, only one part of which is to drive to the locations in question. No sensible business would pay for self-driving trucks and then hire a worker just to ride along until they arrived at a delivery point unless the cost of such trucks became negligible. Gittleman and Monaco (2019) calculate that if autonomous trucks do arrive, the job losses associated with them are roughly one-tenth of what popular accounts are claiming because of the above-mentioned caveats.

Remus and Levy (2016) examine how IT and data-science technologies are affecting the practice of law. This is relevant because the ability to search cases and build legal

arguments can now be done electronically. They conclude that these technologies are not eliminating lawyers – they are simply automating one research-related task, allowing lawyers to focus more time on others. As an example, consider situations where IT simply provides new information used in decisions. As noted above, machine-learning algorithms that read x-rays to look for tumours or interpret other medical tests are not eliminating the doctors who make the diagnosis about a patient. They provide a new and important set of information that is combined with other information – patient histories, blood and genetic tests, and so forth – that doctors use to make diagnoses. It is possible to imagine a future where the entire judgment process is taken over by robots, but that vision is so far away at this point that we are simply projecting it. Autor (2015) also notes that even when new technologies do eliminate tasks, and possibly jobs, the changes take place quite gradually.

Bresnahan and Yi (2016) offer the most sweeping refutation of the notion that new IT will eliminate jobs by reminding us that IT and technology generally alter products and services in ways that give customers additional benefits and features rather than simply automating existing features. They are typically not producing the exact same product or service. As a result, tasks are not necessarily eliminated. The technology itself creates new products and services or aspects of existing services that create new tasks. One such example is the now common experience of shopping online where the website suggests other products and services the shopper might purchase. Some of those products and services may require connection to an employee. Online travel bookings may lead to recommendations for insurance purchases or requests for advice on health issues associated with travel, such as vaccinations. In that case, the new technology has created new services that did not previously exist and

new tasks for humans, thereby increasing the demand for human labour.

We also know that many tasks that appear to be done by IT actually involve workers behind the scenes. Gray and Suri (2019) document an entire workforce that has been created to support – unseen – tasks associated with doing business

on the internet, such as matching individuals' images to their security photos or editing social media content. No doubt at some point those tasks might become automated, but at present it is cheaper and easier to have them done by people. (The jobs are low wage and performed by arms-length contractors so we should not imply that good jobs have been created.)

8. Robotics and automation

Robotics – the field associated with robots – is the arena where we might expect to see the greatest effects on jobs. It seems quite difficult to come up with an exact definition of a robot, but it is clear that it relates to the application of computer-science techniques to tasks that mimic human behaviour, typically involving the physical world. What differentiates robots from machine tools is that robots have some autonomy: their programming allows them to adapt or adjust to change how it responds to circumstances. A metal press may be a sophisticated and expensive tool that increases labour productivity but it is not a robot. If we add computer programming to it so that it can adapt its performance to the differences it perceives in the metal coming into contact with it, then it may well be. Similarly, 'chat bots' that answer questions asked by individuals in conversation form are typically seen as a type of robot even though they do not engage with the physical world. Although the ability to process natural language in the form of human voices is impressive, their ability to adapt – which is central to robotics – rather than simply respond to an array of questions is quite limited.

Because robots are a specific application of IT to human tasks, we might expect their use to be particularly associated with changes in jobs. However, as with other forms of IT, the ability to take on individual tasks does not

necessarily correspond to a complete job. Like the more general aspects of IT noted above, the robotic industry appears to have shifted its focus from efforts to take over complete jobs to efforts to assist workers in jobs by taking over individual tasks, a much simpler outcome than attempting to take over all the tasks an individual has to perform. In this context, it is useful to note that the set of tasks assembled to create jobs that people do is based on both the logic of what humans can do as well as what organisations need. That logic is not the same as what machines and IT can do, so the notion that IT will somehow will neatly map on to existing jobs is mistaken.

Assessing the possible effects of robots on jobs is essentially the same exercise as assessing the effects of IT in general on jobs. There have been some specific efforts to examine investments in robots *per se*, with Acemoglu and Restrepo (2017) attracting the most attention with their study on spending on robotics showing a negative relationship with regional employment. As Mishel and Bivens (2017) point out, such results do not hold for automation other than robots which had a positive relationship with employment. Graetz and Michaels (2018) use evidence across 17 countries and find that a greater use of robots did not have a significant negative effect on employment. Dixon, Hong and Wu (2019) conduct one of the very few studies at the firm

level, using data from Canada on computer use matched to data on firm practices and outcomes. They found that greater computer use was associated with greater employment growth but a reduction in managerial employment as the introduction of robots appears to lead to changes in work organisation. Borjas and Freeman (2019) compare the effects of the introduction of industrial robots (i.e. larger machines and associated with substituting routine labour tasks) vs. immigrants in US manufacturing industries and conclude that the introduction of these robots is associated with a far greater reduction in total employment than the increase in immigrants, as much as two to three workers for each industrial robot.

To summarise, the evidence is mixed. The studies focus on manufacturing *per se* and would not necessarily capture *employment* effects elsewhere, where increased productivity and robotic sales and service may generate jobs in other contexts. Given that, it is surprising that the studies do not find negative effects on employment. The fact that as many find positive as negative effects leads to the conclusion that, as yet, clear evidence of negative employment effects cannot be seen.

The term ‘automation’ has surfaced recently in discussions about the potential effect of AI on jobs, presumably related to the robotics idea of applications specifically designed to replace workers in jobs. In the United States, the discussion on automation first came to the fore while trying to explain the slow growth of employment in US manufacturing after the Great Depression. The fact that productivity appeared to have jumped in manufacturing was seen as consistent with the possibility that IT had ramped up productivity there. As a result, the claim was that investments in technology held down jobs in manufacturing (see, e.g. Perry, 2017).

The problem with this argument is that closer inspection suggested that it was just not true. The apparent jump in productivity in US manufacturing was attributable in part to changes in what counts as manufacturing: companies like Caterpillar that manufacture expensive heavy equipment have also moved into services – repairing and financing equipment. The income from those service operations has been counted towards manufacturing because the company itself is a manufacturing company. To the extent that the sharp increase in manufacturing was real, it seems attributable largely to one industry – computer manufacturing – and that has not continued.

Houseman (2018) explains these developments and notes there is little support for the idea that increasing productivity was eliminating manufacturing jobs. Furthermore, the jump in productivity in the computer industry was not because of improvements in labour productivity of the kind that is evident in typical industries – i.e. fewer workers required to build the same computer or less labour input in the construction of a computer. It is because changes in computer design, especially in computer chips, make the same computer considerably more valuable when productivity is measured in terms of revenue per employee.

A different kind of argument about IT and productivity surfaced in popular discussion around the publication of Robert Gordon’s (2016) contemporary history of economic growth in the US and what it suggests about the future. The history itself is not controversial although surprising to non-experts: productivity growth in the United States hit its contemporary peak in the 1930s as machine-age innovations were adapted to more everyday uses. Since then, productivity growth and the technological change that drives much of it at least have declined, despite repeated claims in the business and policy world that we are always living in a time of unprecedented change.

The part of the argument that generated controversy is Gordon's assertion that, at least in the foreseeable future, there is little evidence of technological changes that will drive faster rates of productivity and economic growth. This argument is essentially a forecast based on how growth came about in the past and looking at the current state of play.

This forecast is quite pessimistic and not particularly popular with the public, although others have made similar claims. Summers (2015), for example, coined the term 'secular stagnation' to describe the low current growth rates, in his view driven by policy mistakes. Other

economists are more optimistic about future growth, including the role that new technology might play (see Teulings and Baldwin, 2014, for these debates). Brynjolffson presents a counter view from his popular writings, that paradigm-breaking IT developments which do not follow the usual rules for growth are on the horizon³.

This discussion about the future of growth might be described as two views talking past each other: Gordon and others saying that current evidence leads to a pessimistic view of future growth; the sceptics saying, beyond what we can see with our current approach, growth will return and may be considerable.

9. Looking past empirical evidence on IT effects

Recent studies by Autor (2018) and by Acemoglu and Restrepo (2019) have articulated in more formal terms the traditional explanation about why improvements in technology and labour-saving techniques do not lead to fewer jobs: productivity increases fuel demand in the economy as a whole, which in turn creates more jobs, albeit typically in other areas than where the initial productivity improvements take place. There are many paths through which the connection between productivity improvements and demand can take place⁴.

When we review the empirical evidence from studies of IT use and jobs, there is no support for the view articulated by Brynjolffson and McAfee (2011; 2014) that IT and associated AI advances contributed to lower job growth. At least in the United States, during the Great Recession the slack labour market that gave

support to such an argument has turned around now and undercut it. The more complex argument that IT use has led to automation of the most routine jobs and expansion of more sophisticated jobs has greater face value, but empirical evidence for it is at best mixed, and there are several studies with results that directly contradict it.

That leaves one more set of arguments where the usual forecasts are left behind. Here the idea is that something is coming in IT and related AI developments that will be different in its effects on jobs than anything we have seen so far. It is not just the technologies themselves that will be different, but how they will interact with jobs will also be different. As noted above, these are not forecasts because they claim explicitly that the future will not look like the past. As such, projections are

³ This 'debate' derives from a TED talk: <https://www.youtube.com/watch?v=ofWKSWqlgI>

⁴ Acemoglu and Restrepo (2019) go further and argue that some productivity improvements may be more labour-saving and less demand-creating than others. They claim that in the face of lower overall productivity growth in recent decades, there has been a shift towards the kind of growth that has less impact on jobs, although they have no direct way of measuring that change and inferring it from lower wage growth which, of course, could have many other causes.

not relevant. Furthermore, the construction of those arguments is inconsistent with what we know about what makes for good predictions, not just in suggesting that prior experience is not a guide to them but also that current examples do not provide a guide.

The examples given by Bresnahan and Yi (2016) show that current data-science tools which generate algorithms for decisions do not necessarily eliminate jobs even in the areas where they are applied. They seem closest to the type of IT innovations that proponents claim will eliminate jobs.

Nordhaus (2015) takes a novel and quite different approach to test directly the claim of a forthcoming, paradigm-breaking advance in IT that will transform business and jobs. He addresses Brynjolfsson and McAfee (2014) explicitly, which is more or less an extension of their 2011 argument. He asks what we would see in the economy if such a development occurred in terms of developments such as the share of capital devoted to IT in the economy. At least in the contemporary economy, he sees little evidence that we are on the way towards such a development.

It might be fair to describe arguments about the future of IT as researchers limiting their interest to analyses of the present, on the

one hand, and ‘expert judgment’ prediction of a future fundamentally different from the past, on the other. It is virtually impossible to refute a claim about something that might happen in the future, especially when the claim itself (effects on jobs) relies on something that has yet to exist (path-breaking IT). There is a joke in the field of forecasting that we are safe in making any claim about the future so long as we do not have to specify when it will come true: we cannot rule out events that may happen in the future, which few people remember, or hold accountable, claims that eventually turn out to be false, and, as noted above, there are short-term benefits in the attention that authors can secure with spectacular claims.

The fact that the current evidence is inconsistent with the general notion that IT innovations will have dramatic effects on jobs does not prove that it is impossible for IT innovations of some kind to ever have such an effect. However, it should considerably lower our estimate as to whether such a scenario is likely. Furthermore, the fact that, as yet, there is no clear evidence for the simple explanations as to the kind of effects that IT is having on the labour force – e.g. eliminating low-wage or mid-level jobs – does not mean that a consensus view will never emerge about such changes. It does mean that acting now on any of those views is not advisable.

10. What to do about an uncertain future

The notion that the future is uncertain is hardly novel, not just with respect to the workplace but related to almost any aspect of human endeavour. It is also wholly unsatisfying not to be able to know with any certainty what to do about the future.

It is common and, in some circles, to still hear people say that we should take our best guess about the future and go with it, even if we know that guess is not very good. In some circumstances that must be right: the building is on fire, there are two different exits, and even delaying the choice until we are more certain is not a smart strategy. But there are

also circumstances where we are not forced to choose, the consequences of being wrong are great, and the consequences of waiting are minimal. If we are climbing a mountain, for example, we will probably wait to get an accurate weather forecast before ascending towards the summit because the cost of waiting is small compared to the cost of making the wrong decision.

With respect to economic and workforce planning, the track record has not been very good at predicting which jobs will be in high demand far into the future. Even if we are reasonably sure that some jobs will decline in importance in the future, retraining programmes are difficult to put in place unless we are also reasonably sure which jobs will be in demand then. A sensible alternative, therefore, is to wait for better information before acting and shortening the time period involved because forecasts are dramatically better the shorter they are.

It is true that government policies often take a long time to set up and execute, and that makes longer-term efforts more attractive. But in that context, our policy attention might be better spent on designing procedures that allow us to respond faster rather than going with longer-term forecasts that have a poor track record.

One approach to faster and more accurate forecasts might be to think about programmes that are executed at the level of the individual employer rather than the economy as a whole. Particularly with respect to changes associated with technology, we know that the spread of new techniques is not instantaneous: businesses with more resources or with strategies better suited to new approaches will go first, while others may never adopt the changes because of their unique cost structures or business approaches. Estimating what will happen to jobs in a given organisation two years on is

far easier and more accurate than estimating what will happen to jobs in the economy as a whole because at least some of the factors that drive outcomes in a given organisation are known and indeed determined by decisions made within that organisation.

Furthermore, if we believe that IT-related technologies may eliminate jobs, intervening when those developments actually do so – within individual employers – is a far better use of resources than putting in place economy-wide programmes that may only be used by a small group of employees at any specific time. We also know that where individuals must transition from one job to another, the easiest way to make those transitions is within the same organisation where their organisation-specific skills remain relevant. Retraining policies that operate within individual employers may also make sense for that reason.

Another general approach to addressing the problem of uncertainty begins with the recognition that even good forecasting models simply tell us the most likely outcome, or in the words of modellers, the ‘point estimate’ of the outcome in question. In most cases, the most probable outcome may not be all that likely, so it is important to know what the second most likely outcome is, as well as the third. Sometimes the second and third outcomes are similar in their implications, in which case it is safer to bet on them than on the most likely outcome. Scenario planning is one technique used to address these situations. Simulations are another, where we have a forecasting model and we change the assumptions or the values of the variables to see what happens.

Once we have a better sense of the outlines of the uncertainty we face, a reasonable approach involves hedging our bets. The world of finance has formalised this practice in the form of options, and the world of management has done

something similar with the idea of ‘real options’, placing bets to hedge against real phenomenon. For example, the probability might be extremely high that there will not be a pandemic, but the consequences if it does happen are high enough that we might at least put plans in place to deal with it should it happen. If it turns out that evidence of dramatic IT-related job changes becomes stronger, it may make sense to place some bets about it occurring even if the odds are still small that those changes will occur. An example of such a bet might be more detailed and fine-grained monitoring of how IT is being used in the workplace.

Even if we were to believe that new IT technologies, whatever they may be, are unlike any we have seen before, that would not suggest that the process through which any such technologies will be introduced is without precedent. The introduction of electricity, for example, was a path-breaking and ‘disruptive’ technology with little precedent. We learned a great deal over time about why it took so long to spread and what determined its advance. If we look at manufacturing, where technological change has been most obvious and studied, we know that its introduction rarely has uniform effects everywhere. In the 1970s, the term ‘productivity bargaining’ was used to describe an approach which began in the UK whereby unions and management negotiated over the terms on which new technology and other productivity-improving approaches would be introduced that would protect as many current jobs as possible and share some of the benefits

of cost savings with employees (e.g. McKersie and Hunter, 1973). A simple accommodation was to let labour-saving play out through attrition and buy-outs rather than mass layoffs.

Arguably, the first ‘robotics’ wave in manufacturing was the introduction of numerically controlled machines, taking over at least some of the most important tasks of machinists. Here, organisations faced a choice as to whether to get rid of their machinists who had performed those tasks, replacing them with engineers proficient in computer programming, or to retrain their existing machinists to take over the programming tasks. Productivity was actually higher in the latter case (Kelley, 1994). The former approach is massively disruptive for employees; the latter much less so (see Keefe, 1991 for an assessment of overall effects on jobs), and employers had considerable discretion as to which one to choose. The policy approach learned from that is first that these two options have very different implications for society and for employees and second that it would have been possible to shape the choices.

The assertion that we should initiate massive retraining programmes now on the chance that new IT innovations will be massively disruptive is not the only option, even if it was feasible to do, nor even the best given what we know first about the lack of evidence for such a disruption and second about how technological change actually plays out. Fortunately, there are better options.

11. References

- Acemoglu, D. and Restrepo, P. (2017), *Robots and Jobs: Evidence from US Labor Markets*, NBER Working Paper No. 23285.
- Acemoglu, D. and Pascual Restrepo, (2019), Automation and New Tasks: How Technology Displaces and Reinstates Labor, NBER Working Paper No. 25684.
- Allen, J. and de Grip, A. (2012), Does skill obsolescence increase the risk of employment loss? *Applied Economics*, 44, 3237-3245.
- Arntz, M., Gregory, T. and Zierahn, U. (2016), The risk of automation for jobs in OECD countries: a comparative analysis, OECD Social, Employment and Migration Papers No. 189. Paris: Organisation for Economic Co-operation and Development Publishing.
- Aum, S., Lee, S.Y.T. and Shin, Y. (2017), Industrial and occupational employment changes during the Great Recession, *Federal Reserve Bank of St. Louis Review* 99 (4), 307-317.
- Automation Commissions Report, (1965), *Monthly Labor Review*, March 1966, 274-277.
- Automation', *Journal of Economic Perspectives*, 29(3): 3-30.
- Autor, D., Katz, L. and Krueger, A. (1998), Computing inequality: Have computers changed the labor market? *Quarterly Journal of Economics*, 113, 4 (Nov):1169-1213.
- Autor, D.H. (2015), 'Why Are There Still So Many Jobs? The History and Future of Workplace Change: An Empirical Exploration', *Quarterly Journal of Economics*, 116(4).
- Autor, D.H., Katz L.F. and Kearney, M.S. (2008), Trends in US Wage Inequality: Reassessing the Revisionists, *Review of Economics and Statistics*, 90(2): 300-323.
- Bárány, Z.L. and Siegel, C. (2018), 'Job Polarization and Structural Change', *American Economic Journal: Macroeconomics* 10: 57-89.
- Bessen, J.E. How Computer Automation Affects Occupations: Technology, Jobs, and Skills (3 October 2016), Boston Univ. School of Law, Law and Economics Research Paper No. 15-49. Available at SSRN: <https://ssrn.com/abstract=2690435> or <http://dx.doi.org/10.2139/ssrn.2690435>
- Bix, A.S. (2000), *Inventing Ourselves Out of Jobs? America's Debate over Technological Unemployment, 1929-1981*, Baltimore: Johns Hopkins University Press.
- Borjas, G.J. and Freeman, R.B. (2019), From Immigrants to Robots: The Changing Locus of Substitutes for Workers, NBER Working Paper 25438.
- Borland, J. and Coelli, M. (2017), Are Robots Taking Our Jobs? *The Australian Economic Review*, 50(4): 377-97.
- Bresnahan, T.F. (1999), Computerization and wage dispersion: An analytical reinterpretation, *Journal of the Royal Economic Society*, 109, 456, F390-F415.
- Autor, D.H., Levy, F. and Richard, J. Murnane. (2003), 'The Skill Content of Recent Technological

- Bresnahan, T.F., Brynjolfsson, E. and Hitt, M.L. (2002), Information technology, workplace organization and the demand for skilled labor: Firm level evidence, *Quarterly Journal of Economics*, 117(1): 339-376.
- Brynjolfsson, E. and 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. and McAfee, A. (2014), *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. WW Norton & Company, New York.
- Brynjolfsson, E., Mitchell, T. and Rock, D. (2018), What can machines learn and what does it mean for occupations and the economy? AEA Papers and Proceedings 108: 43-47.
- Dedrick, J., Gurbaxani, V. and Kraemer, K.L. (2003), Information Technology and Economic Performance: a Critical Review of the Empirical Evidence. *ACM Computing Surveys*, 35(1): 1-28.
- DiNardo, J.E. and Pischke, J.-S. (1997), The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too? *Quarterly Journal of Economics*, 112(1): 291-303.
- Dixon, J., Hong, B. and Wu, L. (2019), The Employment Consequences of Robots: Firm-Level Evidence: available at <https://ssrn.com/abstract=3422581> or <http://dx.doi.org/10.2139/ssrn.3422581>
- European Automobile Manufacturers Association, (2017), *Managing the Transition to Driverless Road Freight Transport*. 31 May: <https://www.acea.be/press-releases/article/driverless-trucks-new-report-maps-out-global-action-on-driver-jobs-and-lega>
- Fairlie, R.W. and Bahr, P.R. (2018), The Effects of Computers and Acquired Skills on Earnings, Employment, and College Enrollment: Evidence from a Field Experiment and California UI Earnings Records, *Economics of Education Review*, 63: 51-63.
- Fano, E. (1991), a 'Wastage of Men': Technological Progress and Unemployment in the United States, *Technology and Culture*, 32(2), Part 1: 264-292.
- Frey, C.B. and M.A. Osborne. (2017), The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114: 254-280.
- Funk, J. (2017), Assessing public forecasts to encourage accountability: The case of MIT's Technology Review, PLOS One: <https://doi.org/10.1371/journal.pone.0183038>
- Gittleman, M. and Monaco, K. (2019), Truck-Driving Jobs: Are They Headed for Rapid Elimination? *ILR Review*, forthcoming.
- Gordon, R. (2016), *The Rise and Fall of American Growth*, Princeton: Princeton University Press.
- Graetz, G. and Michaels, G. (2018), Robots at Work, *Review of Economics and Statistics*, 100(5): 753-768.
- Gray, M.L. and Siddharth, S. (2019), *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Boston: Houghton Mifflin Harcourt.
- Green, F., Felstead, A., Gallie, D. and Zhou, Y. 2001, Computers and Pay, *National Institute Economic Review*, 63-75.

- Harris, K., Kimson, A. and Schwedel, A. (2018), Labor 2030: The Collision of Demographics, Automation, and Inequality, Bain & Company: <https://www.bain.com/insights/labor-2030-the-collision-of-demographics-automation-and-inequality>
- Keefe, J.H. (1991), Numerically Controlled Machine Tools and Worker Skills. *ILR Review*, April 91, Vol. 44, Issue 3, p. 503-519.
- Kelley, M.R. (1994), Productivity and Information Technology: The Elusive Connection, *Management Science*, 40:1: 1406-1425.
- Kristal, T. (2013), The Capitalist Machine: Computerization, Workers' Power, and the Decline in Labor's Share within U.S. Industries, *American Sociological Review*, 78(3) 361-389.
- Krueger, A.B. (1993), How Have Computers Changed the Wage Structure: Evidence from Microdata, 1984-1989, *Quarterly Journal of Economics*, 108(1):33-61.
- Liu, Y. and Grusky, D.B. (2013), The payoff to skill in the third industrial revolution, *American Journal of Sociology*, 118: 1330-74.
- McKersie, R.B. and Hunter, L.C. (1973), Pay, Productivity, and Collective Bargaining, London: Macmillan.
- Mishel, L. and Bivens, J. (2017), The Zombie Robot Argument Lurches On, Washington, D.C.: Economic Policy Institute.
- Mishel, L.R. and Bernstein, J. (1994), *Is the Technology Black Box Empty? An Empirical Examination of the Impact of Technology on Wage Inequality and the Employment Structure*, Washington, D.C.: Economic Policy Institute.
- Nordhaus, W.D. (2015), Are We Approaching an Economic Singularity? Information Technology and the Future of Economic Growth, NBER Working Paper No. 21547.
- OECD Social, Employment and Migration Working Papers, (2018), Automation, skills use and training: <https://www.oecd-ilibrary.org/docserver/2e-2f4eea-en.pdf?expires=1548664102&id=id&ac-name=quest&checksum=AD3C02D0BA97C-28F669CBD3D34516046>
- Perry, M.J. (2017), More Evidence that Technology and Increased Worker Productivity, not Trade, were Behind US Factory Job Losses. Washington, DC: American Enterprise Institute: <http://www.aei.org/publication/more-evidence-that-technology-and-increased-worker-productivity-not-trade-were-the-forces-behind-us-manufacturing-job-losses/>
- Polanyi, K. (1944), *The Great Transformation: The Political and Economic Origins of Our Time*, Boston: Beacon Press.
- Quinn, J.B. (1967), Technological Forecasting, *Harvard Business Review*, March:1-35.
- Remus, D. and Levy, F.S. (2016), Can Robots Be Lawyers? Computers, Lawyers, and the Practice of Law: <https://ssrn.com/abstract=2701092>
- Robert W. and Bahr, P.R. (2018), The Effects of Computers and Acquired Skills on Earnings, Employment, and College Enrollment: Evidence from a Field Experiment and California UI Earnings Records, *Economics of Education Review*, 63: 51-63.

- Schmitt, J., Shierholz, H. and Mishel, L. (2013), Don't Blame the Robots: Assessing the Job Polarization Explanation of Growing Wage Inequality. EPI-CEPR Working Paper: <https://www.epi.org/publication/technology-inequality-dont-blame-the-robots/>
- Schwitzer, J.P. (1965), Computer Effects Upon Managerial Jobs, *Limperatore-07*, 8(3): 233-236.
- Strassman, P. (1997), *The Squandered Computer: Evaluating the Business Alignment of Information Technologies*, New Cannan, CT: Information Economics Press.
- Summers, L. (2015), My Views and the Fed's Views on Secular Stagnation: <http://larrysummers.com/2015/12/22/my-views-and-the-feds-views-on-secular-stagnation/>
- Teulings, C. and Baldwin, R. (eds). (2014), *Secular Stagnation: Facts, Causes, and Cures*, VOC – CEPR: <https://voxeu.org/content/secular-stagnation-facts-causes-and-cures>
- Thomis, M.J. (1970), *The Luddites: Machine-Breaking in Regency England*, Archon Books.
- Thompson, E.P. (1966), *The Making of the English Working Class*, New York: Vintage Books.
- World Economic Forum, (2018), *The Future of Jobs Report*: <https://www.weforum.org/reports/the-future-of-jobs-report-2018>