DRAFT – NOT FOR QUOTATION

Susskind, Daniel, 'Technological Unemployment', forthcoming in Bullock, Justin, et al. (2022), Oxford Handbook of AI Governance.

TECHNOLOGICAL UNEMPLOYMENT[1]

Daniel Susskind[2]
This Draft: 21 September 2021

Abstract: In this chapter, I explore the history of economic thought that focuses on the impact of technological change on the labour market. I show how, until recently, it was hard to study the threat of technological unemployment in the core models in economics. I explain how that has now changed. And I use these new developments in the literature to argue that we now need to take the threat of a decline in labour demand more seriously than before. From the point of view of AI governance, these developments are significant and attention-worthy: work is both a source of income and of meaning, and if technological progress threatens to cause significant disruption to that work, then understanding the nature of this threat – and how we might respond and shape it – is an important task for all policymakers.

Keywords: Technological change; Artificial Intelligence; Unemployment; Inequality; Skill biased technical change; ALM hypothesis

A. Introduction

In the early 1980s, Wassily Leontief, a Nobel-Prize-winning economist, made one of the most provocative claims in economic thought. He feared that what cars and tractors had done to horses at the turn of the 19th century, computers and robots would eventually do to human beings as well – drive more and more of us out of paid work (Leontief 1979; 1983a; 1983b). And today, the world is captured again by Leontief's fear: in the US, 30 per cent of workers now believe their jobs are likely to be replaced by robots and computers in their lifetime; in the UK, the same proportion think it could happen in the next twenty years (Susskind 2020a).

At the time Leontief was writing, his view was not mainstream in the economics profession. And in the decades that followed, it has proven to be very difficult to engage with his concerns in the core models that explore the impact of technology on the labour market. These models, designed to explain certain empirical puzzles that were emerging in labour markets at the time they were conceived, tended to support a far more optimistic view of the impact of technological change on the labour market. There was little room for Leontief's fears. But in the more recent economic literature, there has been a distinct change. New empirical features of the labour market, and striking developments in the capabilities of the

---

[1] Thank you to Anton Korinek for his expert editorial advice, two anonymous referees for their guidance, and Daron Acemoglu, Katya Klinova and all other attendees at the online conference for AI Governance in March 2021 for helpful comments and reflections.
[2] Fellow in Economics, Oxford University; Visiting Professor at King's College London; daniel.susskind@economics.ox.ac.uk.

Susskind, Daniel, 'Technological Unemployment', forthcoming in Bullock, Justin, et al. (2022), Oxford Handbook of AI Governance.

latest technologies, have resulted in the development of a new set of core models. These now do support the possibility of a more pessimistic account of the impact of new technologies.

In this chapter, I explore this history of economic thought that focuses the impact of technological progress on the labour market. I look at the core models that have been developed, describing the succession of empirical puzzles that they were trying to explain and discussing their respective strengths and weaknesses. My aim is to chart the general change of heart that I detect in the economics literature over the past eighty years or so – from a blinkered optimism about the impact of technological change on the labour market, to a creeping pessimism. I conclude that these intellectual developments mean we now can and ought to take the spirit – though not always the substance – of Leontief's fears more seriously than we might have done when he first wrote them down four decades ago.

B. Technology and The Labour Market: A Brief History

The Neoclassical Production Function

An important starting point in the history of economic thought that explores the impact of technological progress on the labour market is the traditional 'neoclassical production function'. Here, the economy is captured by a single function that describes how different factors combine to produce output. And in this setting, new technologies can do three things: they can make labour more productive, so that it is as if there is more labour in the economy; this is 'labour-augmenting' technological change (Harrod 1942). They can make capital more productive, so it is as if there were more capital; this is 'capital-augmenting' technological change (Solow 1956). Finally, they can do both (Hicks 1932).

To economists, this approach is simple and familiar. But it is important to reflect on how it might have constrained subsequent economic thinking. And very broadly, there are two main concerns with this basic set-up. First, it focuses our attention almost exclusively on how technology might 'help' a particular factor – technological progress only appears as an 'augmenting' or 'complementing' force, and there is no way to capture the idea that new technologies might displace workers from a growing range of tasks (Acemoglu and Restrepo 2018a make a similar complaint). In turn, this 'aggregate' approach, which captures the economy in a single mathematical function alone, is opaque on how new technologies affect different factors –the impression is that new technologies simply increase the productivity of the relevant factor at *everything* that it does, leading to more output as a result. Both these features have reappeared in later literature in different ways, as we shall see, and unhelpfully cloud of view of how new technologies might harm workers – now, or in the future.

Skill-biased Technical Change

Until the turn of the 20th century, the 'skill-biased technical change' (SBTC) thesis was the main way that economists thought about the impact of technological progress on the labour market. One simple exposition of the SBTC thesis was captured in what Acemoglu and Autor (2011; 2012) call the "canonical model": this uses the neoclassical production function from before, with two different types of labour, low-skilled and high-skilled, which combine to produce output in a carefully specified way (through a so-called 'constant elasticity of substitution' production function, meaning that a percentage change in the relative wage of

Susskind, Daniel, 'Technological Unemployment', forthcoming in Bullock, Justin, et al. (2022), Oxford Handbook of AI Governance.

the different types of workers always causes a constant percentage change in the relative use of those workers).

The SBTC thesis was designed to explain an empirical puzzle – that during the post-war period the supply of educated workers in many labour markets increased but, at the same time, the 'price' (i.e. the wage) of those workers rose rather than fell relative to those without an education (see, for instance, Acemoglu 2002). The SBTC thesis explained this by arguing that technological change was 'skill-biased' – new technologies (like the digital electronic computer) required skilled people to put them to effective use, and this increased the demand for skilled workers to such an extent that, although their supply had increased, so did their relative price, the so-called 'skill premium' (i.e. the wage of an average college graduate relative to an average high school graduate).

The theoretical and empirical literature on the SBTC thesis is now very large (see, for instance, Katz and Murphy 1992, Bound and Johnson 1992, Goldin and Katz 1998, Bekman, Bound and Machin 1998, Autor, Katz and Krueger 1998; Card and Lemieux 2001; Acemoglu 2002 noted before; and Goldin and Katz 2008). However, although the SBTC thesis is useful in explaining that distinctive rising skill premium, two problems are associated with the wider use of the canonical model in thinking about the impact of technology on work. Both issues are rooted in its reliance on the neoclassical production function. The first is that, in the canonical model, there is no way for technological progress to make either type of worker worse-off in absolute terms – the relative wage of a worker can fall, but not their absolute wage (i.e. technological change is necessarily a 'q-complement' to both types of labour, see Susskind 2020c for a broader discussion of this point). From an empirical point of view, this is a problem because certain workers have already been made worse-off (see, for instance, Autor 2019 on the declining wages of non-college workers). The second problem is that the canonical model is semantically deprived – there is no formal way to understand how or why it is that technological change makes high-skilled workers better at their work. As Autor et al. (2003) put it, "[i]t fails to answer the question of what it is that computers do – or what it is that people do with computers – that causes educated workers to be relatively more in demand".

Routine-Task-Replacing Technical Change

Autor et al. (2003) was originally intended to provide a deeper explanation for the SBTC thesis. But in doing so, it also introduced two innovations into the literature exploring the impact of technology on the labour market: the so-called 'task-based approach', and an influential theory of the capabilities of systems and machines. Together, these have become known as the 'Autor-Levy-Murnane' hypothesis ('the ALM hypothesis'). This line of thinking has been developed further in more recent research (see, for instance, Levy and Murnane 2004, Autor 2013; 2015a,c).

The first innovation in the ALM hypothesis was the distinction between a 'job' and the different 'tasks' that make it up.[3] As noted, one of the problems with neoclassical production

[3] The task-based approach has a rich intellectual history (see Susskind and Susskind 2015 for a detailed discussion). Many classical social theorists, for instance, were interested in the idea, though spoke in terms of the 'division of labour'. Precursors can be found in the economics literature (Zeira 1998, for instance, discusses

functions is that they provide no explanation for why or how technology affects different factors – factors combine to produce output, new technologies allow factors to produce more output, and that is that. We saw this lack of explanation with the canonical model: new technologies allow high-skilled workers to produce more output, but there is no reasoning that details why that is the case or how this process works. By introducing the concept of 'tasks', the ALM hypothesis was able to provide this missing account. The neoclassical production function was broken down into two composite production functions: one that captured how different tasks combine to produce goods (i.e. a 'task-based production for goods') and another that describes how different factors with different capabilities combine to perform those tasks (i.e. a 'factor-based production function for tasks'). When these two functions were combined, one would be able to recover the neoclassical production function once again. But this breakdown allowed for a deeper explanation of how technology affected a particular factor: new technologies helped factors by making them more productive at certain tasks but might harm them by displacing them from others (Susskind 2016).

To see this deeper explanation in practice, think of the personal computer (PC). In the late 1950s and early 1960s, businesses started to make use of mainframe computers. Soon after that, adoption of the PC began: as late as 1980, the US had fewer than one PC per hundred people, but by the turn of the century that figure had risen to more than sixty. And these machines not only became more widespread but more powerful as well: on one estimate, for instance, computational power from 1950 to 2000 increased roughly by a factor of 10 billion (Susskind 2020a). How does this technological progress relate to the changes that were taking place in the labour market at the time? The task-based approach provides a more detailed explanation: these powerful new machines led to an increase in demand for the particular types of *tasks* performed by high-skilled workers who were able to effectively operate these new technologies. This sort of explanation for SBTC, simple and intuitive it may be, could not be provided with a traditional neoclassical production function.

Building on this distinction between 'jobs' and 'tasks', the second innovation was a further distinction between 'routine' and 'non-routine' tasks. The argument was that while systems and machines could readily perform 'routine' tasks, they would struggle to perform 'non-routine' ones. The motivation for this claim relies on a particular conception of how systems and machines operate – that to perform a task, they must follow an explicit set of instructions or rules articulated by a human being. This very closely reflects the 'expert systems' approach to AI, which was particularly influential in the 1980s (Winston 1977, Hayes-Roth et al. 1983 and Russell et al. 1995 provide an overview of this technique). And, according to economists, a task was 'routine' if a human being could articulate how they performed it – following Polanyi (1966), if perform the task relied on 'explicit' rather than 'tacit' knowledge. If this was the case, then it meant it was straightforward to write a set of instructions, based on that human explanation, for a machine to follow – and so these tasks could be automated. Otherwise, the tasks were 'non-routine', the instructions or rules were inexpressible, and so they could not be automated. This distinction between 'routine' and 'non-routine' tasks is now commonplace (see Susskind 2020a, b, c for a fuller history).

---

how factors produce 'intermediate goods' – rather than 'tasks' – which are they combined to produce a unique output). Autor et al. (2003), though, brought it to wider attention in the profession.

Susskind, Daniel, 'Technological Unemployment', forthcoming in Bullock, Justin, et al. (2022), Oxford Handbook of AI Governance.

The ALM hypothesis was powerful because it could explain a different empirical puzzle that emerged in labour markets at the turn of the 20th century. Here, technological progress appeared to harm middling-skilled workers whose pay and share of jobs (as a proportion of total employment) declined in many countries. This phenomenon was known as the 'polarisation' or the 'hollowing-out' of the labour market (see Goos and Manning 2007; Autor, Katz, and Kearney 2006, 2008; Goos, Manning and Salomons 2009, 2014). And the most compelling explanation was provided by the ALM hypothesis – low- and high-skilled jobs were heavily composed of 'non-routine' tasks which could not readily be automated but middling-skilled jobs were made up of 'routine' tasks that were easier to automate. This reasoning could explain the hourglass figure that was emerging in many labour markets.

In time, though, a problem emerged with the ALM hypothesis: many of the tasks that it identified as being 'non-routine', and out of reach of systems and machines, could increasingly be automated. For instance, the tasks of making a medical diagnosis, driving a car, and identifying a bird at a fleeting glimpse were all thought to be out of reach – and yet there are now many systems that can make medical diagnoses, almost all major car manufacturers have driverless car programmes, and there is even an app that can identify a bird at a quick glance (see Susskind and Susskind 2015, and Susskind 2019; 2020a, b, c). Of course, it is reasonable to question whether these tasks have been fully automated: there are certain illnesses these diagnostic systems cannot identify; particular birds that the app cannot recognise; and no driverless car that can yet function without some human attention. But it is important to note the underlying trend: that many 'non-routine' tasks can be automated to an extent that was inconceivable a few decades ago (Susskind 2020a).

At fault was the underlying conception of machine capabilities on which the ALM hypothesis relied. As noted, this was heavily rooted in how systems and machines worked in the 'first wave' of AI that took place in the 1980s – the so-called 'expert systems' approach – but this did not reflect the different approaches that were emerging at the turn of the century. Newer technologies, using advances in processing-power, data storage capabilities, and algorithm design, were no longer reliant on an explicit top-down articulation of instructions from human beings; instead, they were learning from the bottom-up (Susskind and Susskind 2015, Susskind 2020a). Take the system developed at Stanford that, from a photo, can diagnose whether a freckle is cancerous as accurately as twenty-one leading dermatologists (Esteva et al. 2017). How does it function? Not by following an explicit set of rules articulated by a human doctor, but by identifying patterns in a database of 129,450 past clinical cases. Of course, some of the rules it uncovers in that process may be those that a human doctor relies upon but struggles to articulate. But not necessarily so: the system may discover entirely new rules, too. As a result, in this way an increasing range of 'non-routine' activity is within reach of these systems and machines, which operate in a very different way to human beings.[4]

[4] This system is not without critics. One concern, for instance, raised by several of the authors of the Stanford system, is that it was trained on a data set where some images also contained a ruler. If freckles which are considered suspicious are more likely to measured or marked, then the fear is that this system may actually be partly detecting rulers rather than cancers (or, as one study put it, learning "that rulers are malignant", see Narla et al. 2018). How serious is this problem? It is important to note that, in spite of this bias, the original system still appeared to outperform human doctors in identifying cancers in unseen photos of freckles. In any event, as several of the authors of the Stanford system have noted, exploring the biases associated with non-standardised images is an important task when designing these systems (see, for instance, Esteva and Topol 2019).

Susskind, Daniel, 'Technological Unemployment', forthcoming in Bullock, Justin, et al. (2022), Oxford Handbook of AI Governance.

Newer Approaches

More recent approaches in the literature build on the combined shortcomings of the SBTC and the ALM hypothesis. Often, these continue to adopt the task-based approach, helping to resolve some of the analytical opacities that come with using neoclassical production functions. In turn, they are far more agnostic about the capabilities of machines, avoiding the risks of underestimating them as other researchers have done in the past. A good example of this 'task-based, capabilities-agnostic' approach is Acemoglu and Restrepo (2018b).

This latter feature, that these models are often 'capabilities-agnostic', is particularly important when thinking about the future of work. In these newer approaches, the explicit articulations of machine capabilities and functionality that characterised the earlier literature – the SBTC assumption that technology helps skilled-workers, or the ALM hypothesis that 'non-routine' tasks cannot readily be automated – are gone. Instead, there is an implicit recognition that the boundaries to machine capabilities are uncertain and changing. In that spirit, many newer models instead use an endogenously determined cut-off in task space that marks the boundary between activities that can and cannot be automated (for instance, in Acemoglu and Restrepo 2018b, Aghion, Jones, and Jones 2019, and Moll, Rachel, and Restrepo 2021). There is a reliance on the data, rather than theory, to identify where this moveable boundary might lie – for instance, the AI Occupational Impact Measure, the Suitability for Machine Learning Index, and the AI Exposure Score (Acemoglu et al. 2021). And as technological progress unfolds, the cut-off in task space shifts – there is a process of 'task encroachment' at work, where machines gradually, but relentlessly, take on more tasks (see Susskind 2020a and 2020c).

Yet despite this agnosticism, there remains an important sense in which these new approaches still assume that labour is 'more capable' than machines. It is true that the task cut-off that demarcates which tasks can and cannot be automated is no longer fixed, and machines are free to encroach further into the realm of tasks once performed by workers. But an assumption is still commonly made that, if any entirely new tasks are created by technological change, then these tasks will necessarily be those in which workers, rather than machines, have the comparative advantage (see, for instance, Acemoglu and Restrepo 2018a, b, c). I will return to look at this assumption more carefully later in this chapter. But intuitively, it provides a significant new source of optimism when thinking about the future of work: though machines might encroach deeper into the realm of *existing* types of tasks performed by workers, if any *new* types of tasks are created then these will necessarily – by assumption – be those better-suited to workers.

Across these newer models, there is also a more explicit recognition that technological progress, and the process of task encroachment that it generates, has two different effects on the labour market. On the one hand, new technologies can allow capital to *substitute* for workers at particular tasks, reducing labour demand for those activities. (As a shorthand, I will talk about 'technology' substituting for workers.) But at the same time, they can also *complement* workers at other tasks, increasing the demand for labour at those activities. It follows that the overall effect of technological change on the demand for labour depends on the balance between these two forces: if the complementing force is greater than the substituting force, then the overall demand for the work of human beings increases; but if the substituting force is greater, then that overall demand would instead decline. This distinction

Susskind, Daniel, 'Technological Unemployment', forthcoming in Bullock, Justin, et al. (2022), Oxford Handbook of AI Governance.

between the two forces might sound relatively straightforward. Yet until recently, in both informal commentary on the future of work and the formal economic literature, they were either entangled with one another and hard to distinguish – or simply omitted altogether.

Take the substituting force. This is the effect of technological change that captures the popular imagination today. But in the early formal literature, this effect was often absent: in the canonical model, for instance, there "are no explicitly skill replacing technologies" (Acemoglu and Autor 2011). The two types of labour are q-complements, an increase in the quantity or productivity of one necessarily increasing the marginal product of the other. As a result, there is no way for technological progress to make either type of labour worse-off. But now, under the task-based approach, where labour is often modelled in competition with capital to perform different tasks, it is possible to identify this substitution force far more clearly.

Or take the complementing force. This is the force that tends to be neglected in popular debate about the future of work. As Autor (2015a) put it, "journalists and expert commentators overstate the extent of machine substitution for human labor and ignore the strong complementarities that increase productivity, raise earnings, and augment demand for skilled labor". Part of the difficulty here is that while the substituting force is easy to imagine – 'a worker being replaced by a robot', as it is often put – the complementing force is both much less conspicuous and operates in a variety ways. And the more recent literature is often focused on identifying the different ways these helpful countervailing effects might work.

A final difficulty is that, even when these effects are distinguished from one another in the literature, they are often described in very different ways. For instance, Autor (2015a) uses the simple two-part distinction that I adopt: "machines both substitute for and complement labour". But in Acemoglu and Restrepo (2018c), the substituting force is renamed the 'displacement effect' and the complementing force is decomposed into four separate effects: a 'productivity effect', 'capital accumulation', the 'deepening of automation', and a 'reinstatement effect'. Then in Acemoglu and Restrepo (2019), the substituting force reappears as the 'displacement effect', but the complementing force is decomposed into only two different effects: just a 'productivity effect' and a 'reinstating effect'. There is also an additional complication in that the term 'automation' is used in different ways as well: Autor (2015a) uses the term automation in the broadest sense, as in this chapter, to refer to the general use of technology in the workplace, whereas Acemoglu and Restrepo (2018c) uses it in a narrower way, to refer to the "expansion in the set of tasks that can be produced with capital" alone (I call this 'task encroachment'). None of these approaches are incorrect: all are trying to identify the different ways in which same technological progress might increase or decrease the demand for labour but using different labels in carrying out that exercise.

In any event, understanding this intellectual history, and exploring how economists have changed their minds over time, is of practical use in thinking about a future with insufficient demand for the work of human beings. In early models, such an outcome was often not possible: the complementing force was the focus, and the substituting force was neglected. In later models, this outcome was now possible – but still unlikely. Under the ALM hypothesis, for instance, technology no longer complements all types of workers, but instead only complements those workers who can perform 'non-routine' tasks that cannot be automated. In turn, technology substitutes for those workers who perform 'routine' tasks that can be

Susskind, Daniel, 'Technological Unemployment', forthcoming in Bullock, Justin, et al. (2022), Oxford Handbook of AI Governance.

automated. But this approach still encourages skepticism about the idea of insufficient demand for human beings because of this large range of 'non-routine' tasks that cannot be automated. In short then, while the canonical model supported an optimistic view about the threat of automation by making the strong assumption that technology cannot substitute for labour, early task-based models tended to support the same optimistic view but by making the weaker assumption that "the scope for substitution is bounded" (Autor 2015c).

However, the scope for substitution has proven not to be bounded in the way that the task-based literature once expected – 'non-routine' tasks can increasingly be automated. And in response, recent books, papers, reviews, and reports have tried to work out the new limits of machine capabilities, using a variety of approaches (see Susskind 2020a for an overview). But the difficulty with marking out rigid boundaries in this way is that any conclusions one reaches are likely – as with the ALM hypothesis – to swiftly become outdated. A more productive way to think about technological progress is, again, as a process of task encroachment – that machines gradually but relentlessly take on more types of tasks. This involves recognizing that while it is very hard to say what machines might do in the future, it is relatively certain that they will do *more* than they do today. And so, the critical question is whether as new technologies continue this advance, encroaching further into the realm of tasks once performed by human beings, is it possible that the harmful substituting force might eventually overwhelm the helpful complementing force – and the demand for the work of human beings would decline. This is now a question that, at least in theory, it is possible to explore in the newest 'task-based, capability-agnostic' models.

C. Two Types of Technological Unemployment

It was John Maynard Keynes who first popularised the term 'technological unemployment' back in 1930, believing it was a challenge we would "hear a great deal" about in years to come. He thought it would follow from "our discovery of means of economising the use of labour outrunning the pace at which we can find new uses for labour", but beyond those general fears, he offered few details of how it might happen. In what follows I want build on the intellectual history set out before and distinguish more clearly between two types of technological unemployment, two ways that human beings might find themselves without work due to technological progress. These arguments and ideas are drawn from Susskind (2020a).

Frictional Technological Unemployment

The first is 'frictional' technological unemployment. This is not the sort of technological unemployment that Leontief, whose work was mentioned at the outset, had in mind: here, there is sufficient work for human beings to do. In terms of the two forces, the complementing force continues to overpower the substituting force. The problem is that workers who are displaced from their roles by technology are not able to take up those new roles. Though there are many frictions in the labour market that might create this problem, there are three reasons worth highlighting.

The first friction is the skills mismatch – that workers displaced by new technologies may not have the skills to do the work that has to be done elsewhere in the labour market. Historically, this friction has been an important focus. This is natural, given the empirical evidence from labour markets in the 20th century: that workers with more skills tended to have far better outcomes than those with less skills (i.e. the rising 'skill premium'). Economists, for instance, have written of a 'race' between technology and education, implying that people have to learn the right skills to keep up (see, for instance, Goldin and Katz 2009). For that reason, many named the 20th century the "human capital century", capturing the idea that what mattered during that time for an individual's success -- and a country's prosperity -- was their level of education. And politicians at the time saw 'more education' as the most important available response to the challenge of technological change: from US president Bill Clinton hoping in 1996 that "the 13th and 14th years of education [would become] as universal to all Americans as the first 12 are today to US president Barack Obama claiming in 2010 that "in the coming decades, a high school diploma is not going to be enough. Folks need a college degree. They need workforce training. They need a higher education" (Susskind 2020a). All these proposals were attempts to repair the apparent skills mismatch.

But the skills mismatch is not the only important friction. Another is the place-mismatch – that displaced workers may not live in the same geographical location as the work that is created. From a technological point of view, it is interesting to reflect on how, at the start of the Internet era, the rise of new communication technologies led some to predict that these concerns about location would soon no longer matter. This turned out to be a misplaced view – despite "all the hype about the 'death of distance' and the 'flat world', where you live matters more than ever" (Moretti 2013). In many ways, this is to be expected. Technological progress is often associated with geographical rise and fall: in the US, for instance, the success of Silicon Valley and the decline of the Rust Belt are both regional consequences of technological change (Susskind 2020a). A growing body of empirical evidence also supports the growing importance of the place-mismatch: Acemoglu and Restrepo (2020), for instance, show how different local labour markets in the US have vastly different exposure to the use of robots (in large part, because industries differ from place to place); and Autor (2019) argues that the changing nature of jobs in US cities has left non-college urban workers with much less skilled work than they did decades ago (this, he argues, is an important driver of recent declines in non-college wages).

COVID-19 has, in turn, made the importance of place even more apparent. A distinctive feature of the pandemic is that while many white-collar workers have able to retreat to work from home, this has been far less feasible for lower-paid workers in service roles – waiters in restaurants, baristas in coffee shops, retail assistants on the high-street, receptionists in offices – whose roles are far more tightly linked to being physically present in a specific place. One US survey, for instance, found that while 71 per cent of people earning more than $180,000 could work remotely during the pandemic, only 41 per cent of those earning less than $24,000 were able to do so (Susskind 2020a).

A third mismatch is the identity-mismatch. Here, it is not that people do not have the right skills or live in the right place – but they have a conception of themselves that is at odds with the available work, and they are willing to stay unemployed to protect that identify. In South Korea, for instance, half of the unemployed are college graduates – and partly, this may be because they are hesitant to take up the low-paid, low-quality roles that they did not believe their education was preparation to do. Or take adult men of working age in the US, displaced from traditional manufacturing roles by automation. There are some that say these men would rather not work at all than take up so-called 'pink-collar' work – an unfortunate term,

Susskind, Daniel, 'Technological Unemployment', forthcoming in Bullock, Justin, et al. (2022), Oxford Handbook of AI Governance.

but designed to capture the fact that many of the roles that are hardest to automate are disproportionately done by women; in the US, for instance, 92.2 per cent of preschool and kindergarten teachers, 92.2 per cent of nurses, and 82.5 percent of social workers (Susskind 2020a).

Structural Technological Unemployment

The second type of technological unemployment is a structural one. Here, there is simply not enough work for human beings to do, because labour demand declines to the point where workers no longer find it worthwhile to work -- even if those frictions from before are resolved. Before exploring this idea, though, it is important to remember that structural technological unemployment is the most extreme outcome that could follow from a decline in labour demand: there are many other possibilities. Indeed, one problem with the term 'technological unemployment' itself is that it leads us to think that the only way in which the labour market might respond to a decline in the demand for human beings is through the number of 'jobs' that have to do be done. This is not right: the labour market might adjust in several ways to a decline in demand, not only through the *number* of jobs, but also the *nature* of those jobs – their pay and quality, for instance. To those who find the latter an entirely reasonable proposition but the former hard to fathom, I would encourage them to view the two outcomes as sitting on a continuum of possible consequences, and to recognise that if the demand for the work of human beings were to fall substantially, it is plausible that wages might fall so low for certain people that they prefer not to work at all.

The idea of structural technological unemployment is often dismissed on the grounds that that, ever since modern economic growth began about three centuries ago, people have suffered from bursts of concern about new technologies displacing them from their work – and yet, time and again, they have been wrong. There is thus little evidence, so the sceptical line continues, to support the general fear that technological progress would create large numbers of permanently unemployed workers. And it is true that human beings have been displaced but, time and again, they have been able to take up work elsewhere in the economy. In retrospect, the reason for this repeated mistake can be expressed in terms of the two fundamental forces identified before – over time, in predicting the impact of technological change, people have systematically overestimated the substituting force and underestimated the complementing force. As a result, contrary to their fears, there has always been sufficient demand for the work of human beings.

It follows that, in thinking about the future of work, the key question is whether that favourable balance between those two forces is likely to continue. In my view, an important threat to this balance is the process of task encroachment. And the great merit of the emerging 'task-based, capabilities-agnostic' models is that they allow us to explore the impact of this process in a more formal way; as Acemoglu and Restrepo (2018a, c) put it, they allow us to consider the consequences of "the use of machines to substitute for human labour in a widening range of tasks". In what follows, I want to explain why we should take the fear that the balance between these two effects may shift in the future entirely seriously.

In the early literature, as noted, these two effects were either conflated or the substituting force was omitted altogether (i.e. the canonical model). In later literature, these effects were disentangled but the substituting force was still tightly constrained by the assumption that there were firm boundaries to the capabilities of machines (i.e. the ALM hypothesis). In more recent literature, that constraint on the substituting force has been removed (i.e. the newer task-based, capabilities-agnostic models). But there is, in my view, still a residual problem – the process of task encroachment is still considered almost exclusively through its role in strengthening the substituting force and not through how it may weaken the countervailing complementing force as well. Put another way, the process of task encroachment is troubling not only because it widens the range of tasks in which new technologies can substitute for labour but, at the same time, all else held constant, it also narrows the range of tasks in which capital might then complement labour.

The Complementing Force, Weakened

To see how the complementing force might be weakened by the process of task encroachment, consider some of the traditional channels identified in the literature through which technological progress might complement human beings. As noted, the challenge in identifying these challenges is that economists, over the decades, have broken down this complementing force in different ways. Nevertheless, it is still possible to distinguish three general ways in which this helpful force is thought to work: as a 'productivity effect', a 'bigger-pie' effect, and a 'changing pie' effect (Susskind 2020a).

To begin with, new technologies might complement workers *directly* by making them more productive at certain tasks. I call this the 'productivity effect'.[5] When commentators on the future of work describe how new technologies might 'help', 'augment', 'enhance', or 'empower' workers, making them more effective at the activities that they do, this is the effect they tend to have in mind. In a traditional economic set-up, this effect would be analogous to factor-augmenting technological progress when using a neoclassical production function. However, the task-based approach allows us to be far more explicit about what exactly workers are becoming more productive at doing – namely, unautomated *tasks*. And so, new technologies might displace people from certain tasks, but they can also make workers more productive at tasks that have not been automated. When those improvements in productivity are passed on to consumers through lower prices or higher-quality goods, that may increase the demand for the work of human beings.[6]

Here, though, is the problem with this effect: future technologies will surely make some human beings more productive at certain tasks, but this will only raise the demand for their efforts if they remain able to perform those tasks more efficiently than machines (i.e. at a lower unit cost). When that is no longer the case, the relative productivity of human beings becomes irrelevant – machines will simply perform the task instead. Take a traditional craft like candle making (Susskind 2020a). Human beings were once best placed to perform that task – but that ceased to be the case some time ago. The productivity of a present-day tallow-chandler might interest some enthusiasts, but from an economic point of view it is irrelevant – that task is now done by a machine. As the process of task encroachment continues, human capabilities will become irrelevant in this manner for more tasks. In this way, the process not

[5] Acemoglu and Restrepo (2018c) and Acemoglu and Restrepo (2019) also talk of 'productivity effect' in their framework.
[6] This is analogous to the idea in Autor (2015c), that "workers are more likely to benefit directly from automation if they supply tasks that are complemented by automation".

Susskind, Daniel, 'Technological Unemployment', forthcoming in Bullock, Justin, et al. (2022), Oxford Handbook of AI Governance.

only increases the strength of the substituting force, but it narrows the range of tasks in which the complementing force might directly help human beings through this productivity effect.

But technological progress can also complement human beings *indirectly*. If we think of the economy as a pie, for example, another channel that is widely identified in the literature is the 'bigger pie effect'.[7] Here, as productivity improves, incomes grow, the demand for goods and services rises, and so demand also increases for all that the tasks that are needed to produce them. In a similar spirit to before, new technologies might displace workers from certain tasks, but a growing demand for unautomated tasks elsewhere in the labour market could provide them with work instead. Consider Autor (2015b):

> "I think that people are extremely unduly pessimistic ... we neglect the fact that as we create wealth, which we certainly do through productivity improvements, we create more consumption. People want more experiences; they want more goods and services. Consequently, as people get wealthier, they tend to consume more, so that also creates demand."

Or Summers (2013), recounting his experience of being a 1970s MIT graduate student:

> "The stupid people thought that automation was going to make all the jobs go away and there wasn't going to be any work to do. And the smart people understood that when more was produced, there would be more income and therefore there would be more demand."

The concern with this effect, though, is the same as previously: in the future, incomes are likely to be far greater than today, and demand for goods and services will rise along with them, but this will not necessarily lead to an increase in the demand for the work of human beings. Again, this will only be the case if people remain better placed than machines to perform whatever tasks are newly in-demand (i.e. at a lower unit cost). And as the process of task encroachment continues, it becomes more likely that a machine will be better placed to perform those tasks instead. Once again, this process not only widens the range of tasks in which machines substitute for labour, but it narrows the range of tasks in which new technologies might complement workers (here, indirectly).

Or consider another way in which new technologies might indirectly complement human beings – that technological progress not only makes the economic pie bigger, but it also changes the pie: this is the 'changing pie effect'. The British economy, for example, is not only more than a hundred times the size it was three hundred years ago, but its output, and the method of production, has transformed: agriculture now only employs 380,000 people compared to 3.2 million in 1860; manufacturing only employs 40 percent of the number of workers it did back in 1948 (Susskind 2020a). But these trends have not led to vast pools of unemployed people because the economy is no longer made up of farms and factories alone. As Autor (2015c) puts it so well, there was little chance that "farmers at the turn of the twentieth century could foresee that one hundred years later, healthcare, finance, information

[7] If the 'productivity effect' is the partial equilibrium effect of an increase in productivity – workers become more productive at their work, and their wages might rise as a result – the 'bigger pie effect' can be thought of as a general equilibrium effect of an increase in productivity.

technology, consumer electrics, hospitality, leisure and entertainment would employ far more workers than agriculture".

Historically, consumer tastes and preferences have been an important driver of these structural transformations – over time, they not only have larger incomes, but they change how they spend those incomes as well. This means new goods and services are demanded. And so, there may also be a demand for displaced workers to performs tasks that have to be done to produce them. This is popular argument: consider Mokyr et al. (2015), for instance, who writes that "the future will surely bring new products that are currently barely imagined, but will be viewed as necessities by the citizens of 2050 or 2080"; or Autor and Dorn (2013), who claim that technological progress will "generate new products and services that raise national income and increase overall demand for labour in the economy".

In the future, it is entirely possible that human beings will have different wants and needs to today – perhaps currently inconceivable ones. Yet this will not necessarily mean a greater demand for the work of human beings as well. That will only be the case, as with the other helpful effects, if human beings remain best placed to perform whatever tasks have to be done to produce those goods and services. But in the future, as task encroachment continues, machines may simply become the more efficient choice instead. Look at newer parts of economic life, and you can catch a glimpse of how this might unfold: In 1964, the most valuable company in the US was AT&T, which had 758,611 employees. But in 2018 it was Apple, with only 132,000; and in 2019, it was Microsoft with 131,000 (Susskind 2020a). Neither company existed back in the 1960s. Many of the goods and services would have been hard to imagine in a pre-Internet era.

A common theme runs through the previous discussion of task encroachment and the complementing force – while it is right to recognise that as our economies grow and change the demand for tasks to produce everything will grow and change as well, it is wrong to think that human beings will necessary be the most efficient choice to perform many of those tasks. I call this the 'superiority assumption' (Susskind 2020a). Economists describe the demand for labour as a 'derived demand', recognising that workers are only demanded in so far as the goods and services that they produce are demanded. But the task-based approach reveals a deeper sense in which it is a derived demand: workers are only demanded in so far as they are best-placed to do whatever tasks must be done to produce those goods and services. While the superiority assumption holds that might be the case. But as it weakens, it may not.

And so, the argument to take structural unemployment seriously follows: "over time, machines continue to become more capable, taking on tasks that once fell to human beings. The harmful substituting force displaces workers in the familiar way. For a time, the helpful complementing force [working in the ways identified] continues to raise the demand for those displaced workers elsewhere. But as task encroachment goes on, and more and more tasks fall to machines, that helpful force is weakened as well." (Susskind 2020a). Eventually, the substituting force overruns the complementing force and the demand for the work of human beings falls away. This does not necessarily lead to a technological big bang after which large numbers of people suddenly find themselves without work. Instead, a far more gradual withering in the demand for labour is possible, as the balance between the two forces tips out of favour of human beings. From an empirical point of view, one might argue that this sort of phenomenon can already be seen in the data: Acemoglu and Restrepo (2020a), for instance, studied the use of industrial robotics in the US from 1990 to 2007 and found a relatively recent case of the substituting force overpowering the complementing force: one more robot per thousand workers, they observed, reduced the employment-to-population ratio in the US economy by 0.2 percentage points and wages by 0.42 percent. In traditional models of the labour market, it was not possible for new technologies of any types to have these harmful aggregate consequences.

To be clear, this reasoning does not rely on the assumption that machines will be able to do everything in the future. In a recent survey, leading computer scientists made the claim that there is a 50 percent chance that machines will outperform human beings at "every task" within forty-five years (Grace et al. 2018). But it is possible that these fears are wrong: tasks might remain that prove impossible to automate, others than are possible but unprofitable to automate, some that remain restricted to human beings for cultural or regulatory reasons, and others that we prefer human beings to do because we value the very fact they are done by a person and not a machine (crafting a fine suit, preparing a delicious meal, caring for one another in ill health and old age). Yet even though machines may not do everything in the future, despite those expert predictions, they will certainly do *more.* And it is this gradual, but relentless, process of task encroachment that is worrisome: human beings forced into a diminishing set of activities, with no economic law to say there must be enough demand for those residual unautomated tasks to keep everyone who wants a job employed at a sufficient wage.

The Creation of New Tasks

In the new 'task-based, capability-agnostic' models, there is a recognition that the complementing force need not be stronger than the substituting force. For instance, in Acemoglu and Restrepo (2018b) a 'horse equilibrium' is identified where workers are immiserated by technological progress. This label is of course a reference to the work of Wassily Leontief, mentioned at the outset, who thought technology would do the same to human beings as it had done in the past to horses. As noted, when Leontief was writing, his view was far from mainstream and very difficult to capture in the core models. But as this more recent work shows, that is no longer the case.

Nevertheless, although Acemoglu and Restrepo (2018b) identify the horse equilibrium as a possible outcome, it is avoided on the balanced growth path in their dynamic model for an intriguing reason – in this model, technological progress is associated with the endogenous creation of "new and more complex" tasks in which labour has the comparative advantage, and so displaced workers can take up these new activities instead. This is another type of helpful changing-pie effect: not only do consumers buy different goods and services, as before, but producers also changed the way in which they make them. Historically, this process appears to have an important source of demand creation (Acemoglu and Restrepo 2018b, c, and 2019). And it is why, Acemoglu and Restrepo (2018b) argue, it is wrong to compare the economic fate of human beings and horses: the former, if displaced, could move on to perform complex new tasks, but the latter would only ever be suited for pulling heavy loads from place to place and could not find anything else to do.

This additional complementing effect might have played an important role in the past. Yet it is not obvious that it is immune from the previous concerns about the effects of task encroachment, either. In Acemoglu and Restrepo (2018b), new tasks are endogenously created because technical change in the model is 'directed' – when human beings are displaced, they become cheaper, and so there is an incentive to create new tasks for them to do to take advantage of those lower costs. But we might then ask why this effect did not also help displaced horses find new roles – why was there not also an incentive to create new tasks for them to do? This argument might sound a little playful, but there is a very serious point: the reason that new tasks were not created for horses to do was because their capabilities had been exhausted relative to machines. No matter how low their relative cost, there was no realm of economic activity in which they could be useful re-employed at a similar scale. At the moment, human capabilities are so remarkable relative to machines that it seems entirely reasonable to assume that we can continue to devise more and more new tasks in which labour has the comparative advantage. But as task encroachment continues, and human beings start to resemble horses when compared with machines for more activities, that assumption looks increasingly questionable.

In the most recent literature, there is a growing recognition that this endogenous process of task creation alone may be insufficient to maintain the demand for the work of human beings. There is, for example, a fear that the labour market might tend to engage in "excessive automation" (see, for instance, Acemoglu and Restrepo 2019, 2020b, and Acemoglu 2021). Rather than passively rely on the endogenous forces identified before, this work involves a call for the state – through taxes and regulation, for instance – to actively strengthen the incentive to develop technologies that complement, rather than substitute, for human beings. This call resembles an increasingly influential movement in computer science, led by Stuart Russell, to incentivise the development of AI that is "provably beneficial" to human beings (see Russell 2019, for instance): again, the belief is that AI research, left undirected, will develop in a way that is harmful to human beings, and ought to be steered in a more positive direction. Putting to one side the question of whether calls to redirect technological change are desirable or feasible, the most important observation for the purposes of this chapter is that these interventions are being discussed at all. This demonstrates that the problem of technological unemployment, where there is not enough demand for the work of human beings to keep everyone is sufficiently well-paid employment, is not simply a new theoretical possibility but one of increasing practical concern, too, that demands our attention.

D. Taking the Idea Seriously

Earlier in this chapter, I explained that a decline in labour demand need not only manifest in a fall in the quantity of available work, but its quality too – its pay, security, status, and so on. This, in part, explains why those who dismiss technological unemployment as a sudden break from economic life today are likely to be mistaken. As argued in Susskind (2020a), it is not a coincidence that concerns about automation are intensifying at the same time as worries about inequality are growing – the two problems are closely related. Today, the labour market is the main way that we share out income in society; for most people, their job is their primary source of income. The growing inequalities that we see in many labour markets show that this approach is already under stress – some get far more for their efforts than others. Technological unemployment is but a more radical version of this same story, but one that ends with some workers earning nothing at all. In turn, the argument that this outcome is only a threat if *most* people find themselves without work is an unhelpful distraction: a world where even 15 to 20 percent of people, for example, find themselves without sufficiently well-paid work would still present serious challenges.

If we do take these ideas seriously, what challenges will we have to confront? From an economic point of view, the main problem is a distributional one. In the future, technological progress is likely to make us collectively more prosperous than ever before. The task will be to find a way to share out that prosperity if our traditional way of doing so, paying people for the work that they do, is less effective than in the past. When John Maynard Keynes introduced the prospect of 'technological unemployment' he simultaneously dismissed it as a threat – within a century, he thought, we would be so collectively prosperous that the traditional economic problem that haunted our ancestors, the "struggle for subsistence", would be resolved and we could all "live wisely and agreeably and well". His prediction was correct – today, global GDP per head is indeed almost large enough to pull everyone out of poverty. But he also made a big mistake – he assumed that the world's prosperity would automatically be enjoyed by everyone. Already, that is far from the case; the economic pie may be far larger than ever before, but most people's slice of it remains wafer-thin (see Susskind 2020a and Stiglitz 2008 for more detail on this argument). And in a world with technological unemployment, this challenge will be even greater.

But technological unemployment will also present us with problems that take us beyond the traditional questions that occupy most economists exploring the impact of technology on work. It is often said, for instance, that work is not simply a source of an income but also of meaning and direction. Understanding this relationship between work and purpose is critical because its nature will necessarily shape the form of any interventions that are required in the future. Suppose, for instance, that for some human beings the link between work and fulfilment is important. For them, a 'job guarantee' might be an appealing intervention, providing them with an income as well as meaningful activity. Alternatively, suppose you look to the almost 70 per cent of workers in the US are either 'not engaged' in or 'actively disengaged' from it, while only 50 percent say they get a sense of identity from their job. For them, a 'basic income' might be more appropriate, providing an income but allowing them to find purpose beyond the traditional labour market. These are unfamiliar and challenging ideas. They raise hard questions – not simply how to pay for such a scheme, a preoccupation of many who are concerned with proposals like this, but how to maintain social solidarity in a world where some people do not make an economic contribution to the collective pot through the work that they do.[8] But if the threat of technological unemployment is a real one, then we must now take these new challenges seriously as well.

E. References

Acemoglu, Daron and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings". In: Handbook of Labour Economics, Volume 4B eds. David Card and Orley Ashenfelter, 1043-1171. Amsterdam: North-Holland.

[8] For that reason, for instance, some might prefer a 'conditional' basic income, where strings are attached to any financial support, rather than a 'universal' basic income, which is given without any conditions. This distinction is explored at greater length in Susskind 2020 and 2021.

Susskind, Daniel, 'Technological Unemployment', forthcoming in Bullock, Justin, et al. (2022), Oxford Handbook of AI Governance.

Acemoglu, Daron and David Autor. 2012. "What Does Human Capital Do? A Review of Goldin and Katz's The Race between Education and Technology". Journal of Economic Literature 50 (2) 462-463.

Acemoglu, Daron, and Pascual Restrepo. 2018a. "Modelling Automation." *AEA Papers and Proceedings*, 108: 48-53.

Acemoglu, Daron and Pascual Restrepo. 2018b. 'The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment', *American Economic Review* 108 (6) 1488-1542.

Acemoglu, Daron and Pascual Restrepo. 2018c. 'Artificial Intelligence, Automation, and Work', in. *The Economics of Artificial Intelligence: An Agenda* (Chicago: University of Chicago Press).

Acemoglu, Daron and Pascual Restrepo, 2019. 'Automation and New Tasks: How Technology Displaces and Reinstates Labour', *Journal of Economic Perspectives.* 128:6.

Acemoglu, Daron and Pascual Restrepo, 2020a. 'Robots and Jobs: Evidence from US Labor Markets', *Journal of Political Economy* 128:6.

Acemoglu, Daron and Pascual Restrepo, 2020b. "The wrong kind of AI? Artificial intelligence and the future of labour demand," *Cambridge Journal of Regions, Economy and Society*, Cambridge Political Economy Society, vol. 13(1), pages 25-35.

Acemoglu, Daron. 2002. "Technical Change, Inequality, and the Labor Market". Jour- nal of Economic Literature 40 (1): 7-72.

Aghion, Philippe, Benjamin Jones, and Charles Jones, 2019, 'Artificial Intelligence and Economic Growth' in *The Economics of Artificial Intelligence: An Agenda* eds. Agrawal, Gans, and Goldfarb.

Autor, David. 2013. "The "Task Approach" to Labor Markets: An Overview", Journal for Labour Market Research 46 (3) 185 - 199.

Autor, David. 2015a. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation". Journal of Economic Perspectives 29 (3) 3-30.

Autor, David. 2015b. "The Limits of The Digital Revolution: Why Our Washing Machines Won't Go To The Moon". An interview with SocialEurope.eu, https://www.socialeu – rope.eu/ 2015/10/the-limits-of-the-digital-revolution-why-our-washing-machines-wont-go-to-the-moon/. October.

Autor, David. 2015c. "Polanyi's Paradox and the Shape of Employment Growth" in 'Re-evaluating Labor Market Dynamics: A Symposium Sponsored by the Federal Reserve Bank of Kansa City, Jackson Hole, Wyoming, August 21-23, 2014'.

Susskind, Daniel, 'Technological Unemployment', forthcoming in Bullock, Justin, et al. (2022), Oxford Handbook of AI Governance.

Autor, David. 2019. 'Work of the Past, Work of the Future', Richard T. Ely Lecture, *AEA Papers and Proceedings* 109: 1-32.

Autor, David and David Dorn. 2013. "Technology Anxiety: Past and Present". Bureau for Employers' Activities, International Labour Office, December.

Autor, David, Lawrence Katz and Melissa Kearney. 2006. "The Polarisation of the U.S. Labor Market". The American Economic Review 96 (2) 189 - 194.

Autor, David, Lawrence Katz and Melissa Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists". The Review of Economics and Statistics 90 (2) 300 - 323.

Autor, David, Lawrence Katz and Alan Krueger. 1998. "Computing Inequality: Have Computers Changed the Labor Market?". The Quarterly Journal of Economics 113 (4) 1169 - 1214.

Autor, David, Frank Levy and Richard Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration". The Quarterly Journal of Economics 118 (4) 1279-1333.

Bekman, Eli, John Bound and Stephen Machin. 1998. "Implications of Skill-Biased Technological Change: International Evidence". The Quarterly Journal of Economics 113 (4) 1245 - 1279.

Bound, John and George Johnson. 1992. "Changes in the Structure of Wages in the 1980's: An Evaluation of Alternative Explanations'. The American Economic Review 82 (3) 371-392.

Card, David and Thomas Lemieux. 2001. "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis". The Quarterly Journal of Economics 116 (2) 705-746.

Esteva, Andre, Brett Kuprel, Roberto A. Novoa, *et al.* 2017. 'Dermatologist-level Classification of Skin Cancer with Deep Neural Networks', *Nature* 542, 115–18.

Esteva, Andrew, and Eric Topol. 2019. 'Can skin cancer diagnosis be transformed by AI?', *The Lancet* 394 10211:1795.

Goldin, Claudia and Lawrence Katz. 1998. "The Origins of Technology-Skill Complementarity". The Quarterly Journal of Economics 113 (3) 693-732.

Goldin, Claudia, and Lawrence Katz. 2008. The Race Between Education and Technology. Cambridge, MA: Harvard University Press.

Goos, Maarten and Alan Manning. 2007. "Lousy and Lovely Jobs: The Rising Polarisation of Work in Britain". The Review of Economics and Statistics 89 (1) 118-133.

Susskind, Daniel, 'Technological Unemployment', forthcoming in Bullock, Justin, et al. (2022), Oxford Handbook of AI Governance.

Goos, Maarten, Alan Manning and Anna Salomons. 2009. "Job Polarisation in Europe". The American Economic Review 99 (2) 58-63.

Goos, Maarten, Alan Manning and Anna Salomons. 2014. "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring". The American Economic Review 104 (8) 2509-2526

Grace, Katja, John Salvatier, Allan Dafoe, et al., 'When Will AI Exceed Human Performance? Evidence from AI Experts', *Journal of Artificial Intelligence Research* 62 (2018), 729–54.

Harden, Blaine. 1982. 'Recession Technology Threaten Workers', *Washington Post,* 26 December.

Harrod, Roy. 1948. Towards a dynamic economics. London, England: Macmillan.

Hayes-Roth, Frederick, Donald Waterman and Douglas Lenat. 1983. Building Expert Systems. MA: Addison-Wesley.

Hicks, John. 1932. The Theory of Wages. London, England: Macmillan.

Katz, Lawrence and Kevin Murphy. 1992. 'Changes in Relative Wages, 163-1987: Supply and Demand Factors', *The Quarterly Journal of Economics,* 107:1 34-78.

Keynes, John Maynard, *Essays in Persuasion* (New York: W. W. Norton, 1963).

Levy, Frank and Richard Murnane 2004. The New Division of Labor. India: New Age International Ltd.

Mokyr, Joel, Chris Vickers, and Nicholas Ziebarth 2015. "The History of Techno- logical Anxiety and the Future of Economic Growth: Is This Time Different?" *Journal of Economic Perspectives* 29 (3): 31-50.

Moll, Benjamin, Rachel Lukasz, and Pascual Restrepo. 2021. 'Uneven Growth: Automation's Impact on Income and Wealth Inequality', NBER Working Paper No. 28440.

Moretti, Enrico. 2013. *The New Geography of Jobs*

Narla, Akhila, Brett Kuprel, Kavita Sarin, Roberto Novoa, and Justin Ko. 2018. 'Automated Classification of Skin Lesions: From Pixels to Practice', *Journal of Investigative Dermatology,* 138:10 pp. 2108-2110.

Solow, Robert. 1956. "A Contribution to the Theory of Economic Growth." The Quarterly Journal of Economics 70 (1): 65-94.

Stiglitz, Joseph. 2008. 'Toward a General Theory of Consumerism: Reflections on Keynes's Economic Possibilities for Our Grandchildren' in Lorenzo Pecchi and Gustavo

Susskind, Daniel, 'Technological Unemployment', forthcoming in Bullock, Justin, et al. (2022), Oxford Handbook of AI Governance.

Piga (eds), *Revisiting Keynes: Economic Possibilities for Our Grandchildren* (Cambridge, MA: MIT Press, 2008).

Summers, Lawrence. 2013. "Economic Possibilities for Our Children". The 2013 Martin Feldstein Lecture. NBER Reporter, No.4.

Susskind, Daniel. 2016. 'Technology and Employment: Tasks, Capabilities, and Tastes', Thesis for DPhil in Economics, Oxford University.

Susskind, Daniel. 2019. 'Re-thinking the capabilities of technology in economics', *Economics Bulletin* 39 (1), 280-288

Susskind, Daniel. 2020a. *A World Without Work: Technology, Automation, and How to Respond* (Allen Lane: London)

Susskind, Daniel. 2020b. 'Work in the Digital Economy', in *Work in the Future* eds. Robert Skidelsky and Nan Craig (Palgrave MacMillian: London)

Susskind, Daniel. 2020c. 'A Model of Task Encroachment in the Labour Market', Oxford University Working Paper.

Susskind, Daniel. 2021. 'A World With Less Work', A Boston Review Forum on Redesigning AI, (May).

Susskind, Daniel, and Richard Susskind. 2015. *The Future of the Professions: How Technology Will Transform the Work of Human Experts* (OUP: Oxford)

Russell, Stuart. 2019. *Machine Compatible* (Allen Lane: London)

Russell, Stuart and Peter Norvig 1995. Artificial Intelligence: A Modern Approach. New Jersey: Pearson Education.

Winston, Patrick. 1977. Artificial Intelligence. Reading: Addison Wesley.