Oxford Review of Economic Policy, 2025, 41, 313–325
https://doi.org/10.1093/oxrep/graf025
Article

Directed technological change: a history and a critical agenda

Daniel Susskind

Gresham College, King's College London, Institute for Ethics in AI and Economics Department, University of Oxford, UK, Stanford Digital Economy Lab, USA, e-mail: danielsusskind@gmail.com

Many thanks to David Vines, Ian Goldin, two anonymous referees, and the attendees at the 27 June 2024 seminar on 'Future of Macroeconomics III' for helpful comments and guidance.

Abstract

In this article, I explore the idea of 'directed technological change'. Although it has its intellectual origins in early twentieth-century economic thought, the idea was largely ignored in the decades that followed. This only changed recently, when contemporary economists adopted and reworked the idea, applying it to some of our great challenges: from climate change to the impact of AI on work. In what follows I explore the history of the idea, set out its modern theoretical form and practical relevance, before closing with important limitations: technical, political, and moral.

Keywords: directed technical change, technological progress, climate change, AI.

JEL codes: B22, O25, O3

I. The history of the analysis of technological change: direction and amount

The possibility that technological progress might have a 'direction'—that it might be biased towards particular factors of production—is present in early formal models. So too is the idea that this direction might be the outcome of deliberate choice, i.e. that technical progress might be 'directed'. Consider the 'neoclassical production function', a cornerstone of macroeconomic modelling (Susskind, 2022). Here, the economy is captured by a single function that describes how different factors combine to produce output:

$$Y = A_3 F(A_2 K, A_1 L) \tag{1}$$

It was recognized early on that technological change can have three different 'directions': it can make labour more productive, known as 'labour-augmenting' technological change (A_1, Harrod, 1948); it can make capital more productive, known as 'capital augmenting' technological change (A_2, Solow, 1956); and it can do both (A_3, Hicks, 1932). Moreover, there was interest in what *determines* this direction. The discussion of what can cause technological change to take a particular form is often said to have begun with John Hicks in the 1930s. (Hicks used the term 'induced technical change' rather than 'directed technological change'.) In *The Theory of Wages*, he argues that:

> A change in the relative prices of the factors of production is itself a spur to invention, and to invention of a particular kind—directed to economising the use of a factor which has become relatively expensive … [W]e need to distinguish two sorts of inventions. We must put on one side those inventions which are the result of a change in the relative price of the factors; let us call these 'induced' inventions. The rest we may call 'autonomous' inventions. (Hicks, 1932, pp. 124–5)

Thus Hicks not only claims that technological progress might have a particular direction—i.e. that it might 'economize in the use' of a particular factor—but he also identifies what might determine that direction—'a change in the relative price of the factors'. Put another way, Hicks articulates for the first time how the *incentives* that individuals face in the market might shape the sorts of technologies that get developed and adopted.

However, despite the power of Hicks' original observation, the idea of directed technological change lay dormant in economic thought during the second half of the twentieth century, as did the idea that this direction might be determined by deliberate choice in response to economic incentives. Instead, the analytical focus was on understanding the aggregate quantity of technical progress that takes place in an economy, rather than its nature: on the *amount* of technical progress (i.e. the magnitude of A), rather than its *direction* (i.e. whether A_1, A_2, or A_3) and the determinants of that direction.

To an extent, this focus in economic thought—towards the amount of technological progress, rather than its direction—is understandable. Robert Solow and Trevor Swan, in a series of influential articles on economic growth published in the late 1950s, had drawn the field's collective attention to the importance of technological progress for sustaining the growth process. It was technological progress, they claimed, rather than the accumulation of capital, that had proved capable of overpowering the process of diminishing returns that, according to the classical economists, had kept humankind stuck in a 'stationary state' (Solow, 1956, 1957; Swan, 1956, 1960).[1] There was thus much work devoted to showing that the level of economic growth was determined by the amount of technological progress (see, for example, Maddison (2001), an attempt to understand growth—and the lack of it—during the last millennia).

But, inspired by Solow and Swan, this work treated all technological progress as exogenous. It thus left a big explanatory hole: where the technological progress came from. And, more than that, what direction did the progress take—i.e. in what way was the progress biased, what determined that bias?

In fact, much of this early work was hampered by the use of a Cobb–Douglas production function. If that functional form is used to underpin a regression designed to identify the effect of technical progress on economic growth analysis of the causes of economic growth it will not be possible to use the data to determine the direction which technological progress has actually taken. That is to say, it will not be possible to distinguish whether the progress is due to an increase in A_1, A_2, or A_3.[2]

During the later part of the twentieth century, economists began to fill the hole left by Solow–Swan framework. Various attempts to endogenize the aggregate process of technological progress followed. To begin with, Lucas (1988) looked to the labour market: people could not only invest in traditional capital, as in Solow–Swan, but also in human capital, and as people gradually became more productive, that could offset the diminishing returns; Romer (1986, 1990, 1992, 1993) generalized Lucas's insight, shifting focus from the world of tangible objects to the intangible world of ideas, and exploiting their peculiar property of non-rivalry in consumption. (After all, what is 'human capital' but ideas stored in the heads of human beings.) Then, at the turn of the century, models followed that saw 'creative destruction', the productivity-improving jostling of rival innovators, as key (Grossman and Helpman, 1991; Aghion and Howitt, 1992).

Although their approaches differed, these economists still shared a common ambition: to provide a deeper explanation for the *amount* of technological progress that might take place. To be clear, there were important exceptions, where economists explored the direction of technological progress and its determinants, as well as the amount of technological progress. There were, for example, direct responses to Hicks (1932): Bloom (1946) argued that the value of the factor price ratio was important in determining the direction of technological progress, not only changes in it; Salter (1966) and Kennedy (1964) argued

---

[1] 'Malthus disclosed a devil', wrote John Maynard Keynes, and 'for half a century all serious writings held that Devil in clear prospect' (Susskind, 2024).

[2] That is simply because the production function takes the form $Y = A_3(A_2K)^{\alpha}\ (A_1L)^{(1-\alpha)}$ in which case the effect on output of an increase in A_2 of a particular amount has the same effect as an increase of A_3 raised to the power α, and similarly the effect on output of an increase in A_1 of a particular amount has the same effect as an increase of A_3 raised to the power $(1–\alpha)$.

that the direction of technological progress could be determined by firms developing technologies for other ends other than economizing on the use of the most expensive factor. In turn, there were debates about the role of relative factor endowments in determining the capital intensity of technological progress in the US relative to the UK during the nineteenth century (for instance, Rothbarth, 1946; Habakkuk, 1962; Temin, 1966; Clarke and Summers, 1980) and more careful studies of the role of factor prices in shaping technological progress (see Acemoglu (2003) on Kennedy (1964), Samuelson (1965), Drandakis and Phelps (1965)).

But by and large, interest in technological progress focused on endogenizing explanations for the amount, rather than the direction, of the process. This is not simply an interesting detail of intellectual history. It matters because some of the most insightful—and influential—stories that contemporary economists have developed about the impact of technological progress appeal to the fact it has a particular direction, that it has a specific bias: 'for many problems' explored by economists, wrote Daron Acemoglu, 'whether technical change is biased towards particular factors is of central importance' (Acemoglu, 2002). And yet, despite that importance, the deeper determinants of those biases tended to be left entirely unexplained, the direction of technological progress left exogenously determined just as the amount of technological progress had been in the past.

To see this omission in practice, consider the 'skill biased technical change' (SBTC) thesis. Until recently, this dominated economic thought about the impact of technological progress on the labour market. (This literature is large: see Bound and Johnson, 1992; Goldin and Katz, 1998, 2008; Bekman *et al.*, 1998; Autor *et al.*, 1998; Card and Lemieux, 2001; Acemoglu, 2002.) This thesis was developed in response to an empirical puzzle: that the wage of college educated workers relative to high-school educated—the so-called 'skill premium'—was rising in many labour markets in the second half of the twentieth century, despite the increase in the relative supply of the former. How to explain this phenomenon? The SBTC thesis argues that technological progress had a particularly direction: specifically, it was biased towards skilled workers. New technologies, like the personal computer and the software that came with it, required skilled workers to put them to effective use, raising the relative demand for their efforts.

Yet a key question remains for the SBTC thesis: *why* was technological progress skill-biased? It need not be that way: indeed, at different moments in economic history technological progress has exhibited different biases. During the Industrial Revolution in Britain, for example, technological progress benefited unskilled rather than skilled workers: new technologies were 'de-skilling', making it easier for less-skilled people to produce the high-quality output that would have required more skilled workers in the past (Susskind, 2020*a*). This explanatory gap did not go unnoticed in the literature. But the attempts to provide deeper explanations for the biases were themselves not deep enough. David Autor *et al.* (2003), for instance, focused on this omission:

> A wealth of quantitative and case-study evidence documents a striking correlation between the adoption of computer-based technologies and the increased use of college-educated labor within detailed industries, within firms, and across plants within industries. This robust correlation is frequently interpreted as evidence of skill-biased technical change. Yet, as critics point out, this interpretation merely labels the correlation without explaining its cause. It 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. (pp. 1279–80)

And in response, the authors introduced the important 'Autor–Levy–Murnane' thesis (ALM for short) to provide an explanation: in short, they argued, new technologies substituted for human workers at 'routine' tasks but complemented them at 'non-routine' tasks—and the latter tended to be performed by high-skilled workers. But despite the explanatory power of the ALM thesis, it too was vulnerable to the authors' original complaint: it failed to answer *why* the sorts of technologies that were developed at the time were biased towards performing 'routine' rather than 'non-routine' tasks at that moment. There was only one sort of technology in the model—'computer capital'—with a price that declined exogenously, and a bias towards performing particular tasks (i.e. 'routine' ones) that was also set exogenously, based on the authors' understanding of how it was that computers at that time must work (i.e. in order to automate a task, a machine had to follow explicit 'rules' that reflected the precise steps that a human being would also have to follow to perform the same task, see Susskind (2016, 2019)). Put differently, it failed to explain the determinants of the direction that technological progress appeared to exhibit.

II. The modern analysis of directed technological change

At the turn of the twenty-first century, the concern with the direction of technological progress was reborn, and new determinants for that direction were identified. The idea of 'induced technological change' was picked up and refined, under the stewardship of Daron Acemoglu (Acemoglu, 1998, 2002, 2007, 2010, and Acemoglu *et al.*, 2012). To begin with, the concept was renamed 'directed technological change', as it is known today. This was not simply an exercise in rebranding: it was a more appropriate label, capturing not only that technological progress could have a particular direction, or bias, but—importantly—that this progress could be *directed,* actively turned towards a specific bias (Susskind, 2024). This latter point is particularly important for thinking about how this concept might inform practical policy-making—an observation I return to in a moment.

More formally, Acemoglu provided missing micro-foundations for the direction of technological progress, showing how it could be determined by choices taken by profit-maximizing firms. Imagine a setting where a firm is deciding to invest to develop new technologies, each of which complements a particular factor, making them more productive. What determines which technologies get developed, and the direction of technological progress as a result? Acemoglu's argument is that it is the relative profitability of those different technologies that matters. In turn, two conflicting forces determine the relative profitability of different technologies: a 'price' effect, which creates an incentive to develop technologies that are used to produce more expensive goods; and a 'market size' effect, which creates an incentive to develop technologies that have a larger market (see Acemoglu, 2002).

This theoretical account helped to clarify ambiguities that had characterized earlier discussions of the direction of technological progress, set out before. To begin with, it addressed the historical disagreements about the determinants of which technologies are adopted: this work showed it was not primarily the factor price ratio or its value, for instance, as earlier explorations of the determinants of the direction of technological progress had suggested, but the relative profitability of the different technologies that are available. In turn, it showed that, while relative factor endowments were important in determining relative profitability, there was an interesting indeterminacy: the 'price' effect strengthened the incentive to develop technologies that economize on the use of scarce factors (for that scarcity is what makes the goods expensive); the 'market size' effect to develop technologies for the abundant factors (for that abundance is what creates a bigger market for the technology).

With the idea of directed technological change more carefully developed, it was then applied in three different ways.

The first was the labour market: to explain the fact that, at different moments in time, technological progress—as noted before—had exhibited biases towards different factors, complementing low-skilled workers in the late eighteenth and early nineteenth centuries and complementing high-skilled workers in the second half of the twentieth century (Acemoglu, 1998, 2002). The SBTC thesis explanation of the latter had been that new technologies exhibited an exogenously-given bias towards skilled workers. But the directed technical change account was different: the increase in supply of skilled workers in the twentieth century *endogenously* created 'a tendency for new technologies to be skill biased' (Acemoglu, 2002). In turn, directed technical change could also explain the fact that technological progress had been biased in the opposite direction at the start of the Industrial Revolution: the development of deskilling technologies 'coincided with a substantial change in relative supplies… a large migration of unskilled workers from villages and Ireland to English cities'. And so, at that time there was an incentive for profit-maximizing firms to develop new technologies that put these unskilled workers to ever-more productive use.

The second domain in which the idea was applied was the environmental setting (see, for instance, Acemoglu *et al.*, 2011, 2012). Here, progress was taken to exhibit a bias towards either 'dirty' or 'clean' technologies. And importantly, the focus was less on explaining the direction of progress in the past—as in the labour market before—and more on shaping the direction in the future: '[a] satisfactory framework for the study of the costs and benefits of different environmental policies', write the authors in the working paper, 'must therefore include at its centerpiece the endogenous response of different types of technologies to proposed policies' (Acemoglu *et al*., 2011). The policy implications were clear: in the absence of intervention, existing incentives would direct technological progress in a 'dirty' direction, but with the right interventions—environmental regulation, carbon taxes, profit taxes, research subsidies—it was possible to change those incentives in such a way that progress could be redirected in a 'clean' direction. What's more, this intervention would only need to be temporary, because once the 'clean' technologies were sufficiently established, new technologies would continue to be developed in that direction.

And the third domain was the international setting, where the idea was used, in a variety of ways, to explain differences in the adoption and development of new technologies across different countries.

One important use is historical: for some, notably Allen (2009*a*, *b*, 2011), directed technological change not only provides a powerful explanation for why the Industrial Revolution began but why it was *British* (rather than, say, German, French, or Flemish). Allen's account begins with the observation that Britain faced a distinctive set of factor prices—wages were 'remarkably high' relative to other countries and energy was 'remarkably cheap'—and that, in turn, created a unique set of economic incentives for British industrialists: to develop technologies that saved on expensive labour and exploited cheap energy. The result was that, in Britain, technological progress was biased towards the labour-saving, energy-intensive innovations that started the Industrial Revolution—the steam engine, spinning jenny, roller spinner, power loom, and so on.

Other uses of the idea in the international setting are more contemporary: for example, directed technological change is used to explain not only why certain countries adopt particular technologies, but also why others do not. For instance, it provides a revealing account for why Japan today is a world-leader in nursing robotics: they have one of the largest elderly populations in the world, and a historical antipathy towards foreign migrants working in their public services, which together creates a strong incentive to develop technologies that can perform healthcare tasks like Paro, a therapeutic robotic seal, and Robear, a machine capable of lifting immobile patients (see Susskind, 2020*a*). But the idea also explains why certain developing countries, for instance, might not adopt existing technologies that have promoted development in other places—as with Britain's rivals in the early moments of the Industrial Revolution, they cannot be 'forced' to adopt technologies if the economic incentives to use them are not right.

III. Policy implications

And so, to summarize. The early literature showed the importance of technological progress in driving growth. However, this process, despite that importance, was left exogenous, entirely unexplained. In time, though, efforts were made to provide an endogenous account of how the *amount* of technological progress was determined. In much the same way, this literature also showed the significance of the bias of technological progress in determining its impact. But again, despite this, the factors affecting these biases were taken as given, left unexplained. And in response, after yet more time, efforts were made to provide an endogenous account of how the *direction* of technological progress was determined—this is directed technical change.

This new account of how the direction of technological progress was determined was important from a technical point of view, as noted. But it is also significant from a practical point of view. The first reason for this significance is the simplest: this work establishes the fact that that technological progress has a direction. All too often, when political leaders talk about technological progress, and the economic growth associated with it, the implicit metaphor they have in mind imagines the economy as a train: policy-makers can try to push down on the lever and increase technological progress or pull back on the lever and try to slow it down, but the rails themselves are fixed, and the direction of travel is taken as given. Take, for instance, Dorling (2021, p. 2), when setting out his worries about growth slowing down:

> Imagine that you have spent your life on a speeding train and you suddenly feel the brakes being applied. You would worry about what was about to happen next. Now imagine that not just you

> but all the people you know… have lived on that very same speeding train, and that the train has been accelerating for virtually all of their lives. An era is ending.

But that is not the right metaphor for thinking about technological progress. A better alternative is a nautical one (Susskind, 2024). Policy-makers are better thought of as being akin to sailors in a boat: they can choose to pull up the sail and go faster or pull it down and go slower, but they can also change their direction of travel. It may be particularly important to change their direction of travel if—suddenly—it appears that they are heading towards some rocks at the edge of an island which they had not known about.

The second practical contribution, alongside showing that technology has a direction, is showing how policy-makers can actively shape that direction of travel. By endogenizing the direction of technological progress and showing how that direction emerges from the profit-maximizing behaviour of individuals, this literature highlighted the importance of *incentives* in shaping the direction of travel. And most importantly, these incentives are not fixed features of economic life that must be taken as given but can themselves be shaped through interventions—changes to taxes and subsides, laws and regulations, and social cultures and norms (Susskind, 2024).

To see the idea of directed technological change in practice, consider its recent application in thinking about the impact of technological progress on the labour market. Broadly, new technologies are thought to have two distinct effects on work. On the one hand, a technology *substitutes* for human workers when it displaces them from particular tasks, reducing the demand for them to perform those activities. This, when it happens, is relatively easy to see. On the other hand, a technology *complements* human workers when it raises the demand for them to perform other unautomated tasks. This latter helpful effect is more difficult to identify than the harmful one, in part because it can work in a variety of different ways. (Susskind (2022) explores the history of this distinction in detail, parsing the different ways it can work.) For the purposes of this paper, though, I proceed with the simple account: that the substituting force reduces the demand for human workers to perform particular tasks, and the complementing force increases the demand for human workers to perform other unautomated tasks. Looked at in this way, it follows that the aggregate impact of technological progress on the demand for human labour depends on the balance between these two rival forces (Susskind, 2016, 2020*a*, 2020*b*, 2020*c*, 2022).

For a long time in the economic literature, the conventional wisdom was that technological progress would have a benign impact on the labour market. Take, for instance, the SBTC thesis from before, which dominated economic thought with respect to the effect of technological progress in the labour market in the second half of the twentieth century. In the so-called 'canonical' model, which provided the formal foundations for the SBTC thesis, it was not possible for new technologies to make any workers worse off, in absolute terms (i.e. real wage falls were infeasible): all workers would be made better off, but some would be made better off than others.

Put more generally in terms of the two fundamental forces, the view was that the helpful complementing force would tend to overwhelm the harmful complementing force: '[j]ournalists and even expert commentators,' wrote David Autor in 2015 when explaining why there are still so many jobs despite the periodic bursts of anxiety about the threat of automation that have taken place since modern economic growth began, 'tend to overstate the extent of machine substitution for human labor and ignore the strong complementarities between automation and labor… [that] augment demand for labour' (Autor, 2015).

However, in recent years, a change of heart has taken place in the economic literature—from blinkered optimism to a creeping pessimism (Susskind, 2019, 2022). In large part, this is driven by remarkable advances in the capabilities of new technologies, and the fact that many tasks which were thought to be out of reach of automation—and, in particular, so-called 'non-routine' tasks like driving a car and making a medical diagnosis—can now be automated. At its most optimistic, the literature is more agnostic about the impact of automation, uncertain about the balance between the substituting and complementing force; at its more pessimistic, there is a recognition that 'excessive automation' may now be under way, that the balance between the two forces is suboptimal from a social point of view (Acemoglu and Restrepo, 2019, 2020; Acemoglu, 2021).

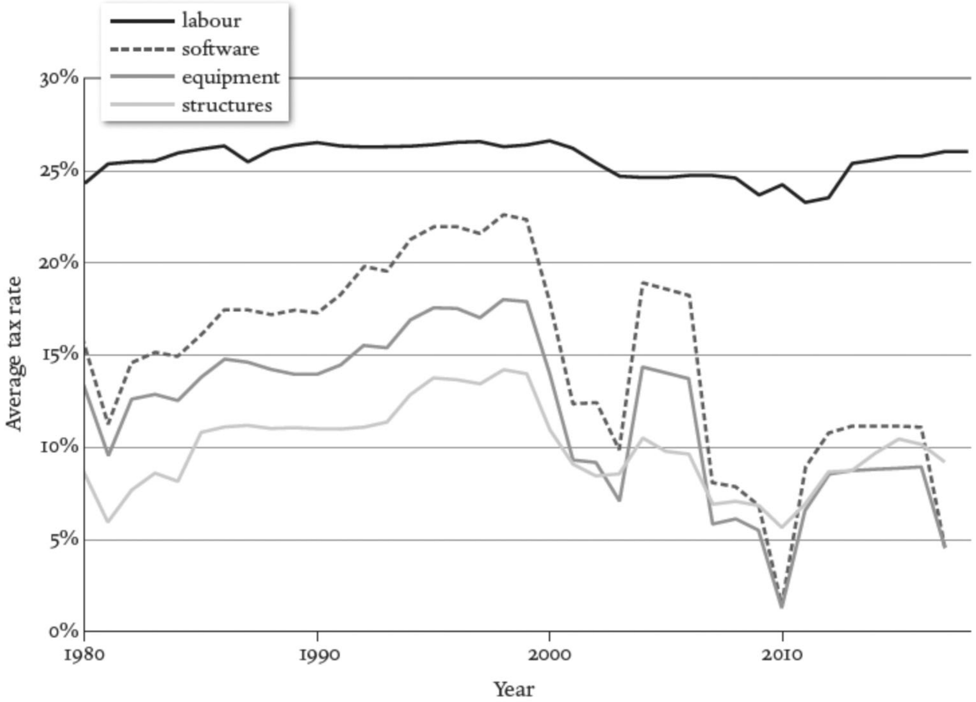

Figure 1: Effective taxes in the US, 1981–2018 (in %). *Source*: Acemoglu *et al.* (2020) in Susskind (2024).

In response, there has been a surge of interest in the possibility of changing the direction of technological progress with respect to the labour market. If the risk is that new technologies might substitute for workers too much, and complement workers too little, then by changing the economic incentives that individuals face in economic markets it might be possible to change the direction of technological progress as well—encouraging the development of technologies that complement, rather than substitute for workers (see, for instance, Acemoglu and Johnson (2023*a*) for a recent statement of these ideas). More specifically, the proposal is that the state can intervene with the toolkit set out before—taxes and subsidies, laws and regulations, social norms and customs—to weaken the incentives that encourage the development of technologies that reduce the demand for human workers, and strengthen those that do the opposite.

To see this in action, consider the first of the three sets of tools—taxes and subsidies. In the US, in every year since 1981, the effective tax rate on human 'labour' has been far higher on 'structures', 'software', or 'equipment' (Acemoglu *et al.*, 2020, in Susskind, 2024). This is shown in Figure 1.

According to certain economists, the practical implication of the trend captured in Figure 1 is that the 'US tax system is biased again labor and as a result generates excessive automation and suboptimally low levels of employment and labor share' (Acemoglu *et al.*, 2020). And so, if this is right, there is an opportunity for the state, by intervening to change these relative tax rates—and, in certain cases, introducing an 'automation tax' on the use of capital to perform tasks in which labour retains the comparative advantage—to improve welfare. (Acemoglu *et al.* (2020) use a task-based model to show that, in theory, the optimal tax rate on capital and labour depends on the inverse supply elasticities of capital and labour as well as any labour market frictions, and then derive the theoretical conditions under which an 'automation tax' is optimal; using 'plausible ranges for the elasticities of the capital and labor supply and estimates of labor market distortions' for the US, they argue that welfare in the US would rise with such a tax.)

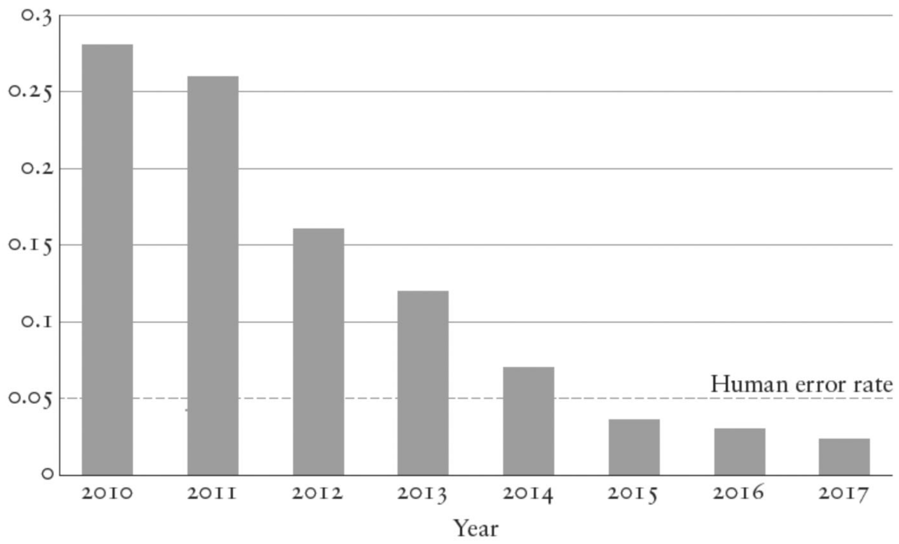

Figure 2: Error rate of the winning system in the ImageNet contest. *Source*: Susskind (2020a).

Or consider another set of tools that policy-makers can use to shape the incentives that people face with respect to the development of new technologies in the labour market: social norms and customs. As others have noted, a distinctive feature of the field of computer science is that the benchmark for success in most domains is human parity (see Acemoglu and Johnson, 2023*a*; Susskind, 2024). That is how technology companies tend to talk about their goals: from DeepMind's ambition to 'solve intelligence', to Facebook's hope of building 'human-level AI'. This is how progress is monitored: from the Stanford AI Index to the Electronic Frontier Foundation's AI Progress Measurement Project, the comparative standard is human performance. And this is how competition in the field is organized. To see this, consider, the ImageNet Contest, an annual event where leading computer scientists competed to build systems that could identify objects in an image more accurately than human beings (Susskind, 2020*a*). Progress in the competition is shown in Figure 2.

There are three significant observations to make about Figure 2. The first is the impressive progress than was made from 2010 in reducing the error rate of the winning system. The second is that, in 2015, the winning system outperformed human beings for the first time—and then improved its performance further in 2016 and 2017. But the third is that the data come to an end in 2017. This is because the competition was stopped. The benchmark was human parity: that was achieved in 2015 and, after some small improvements, the collective attention of the field moved on to other tasks.

Yet one might ask what the consequence of this pervasive human parity benchmark is likely to be on the field—and, more specifically, whether it encourages researchers to develop technologies that substitute for human workers (i.e. dropping below that human error rate in Figure 2) or complement for human workers, pursuing some other non-human benchmark: 'The dominant intellectual paradigm in today's digital tech sector,' write Daron Acemoglu and Simon Johnson, 'is to attain human parity in a vast range of cognitive tasks… This intellectual focus encourages automation [substation] rather than the development of human-complementary technologies' (Acemoglu and Johnson, 2023*b*). But it is possible to change the benchmark for success, and shape the direction of technological progress as well: suppose, for instance, that state-support research bodies commit to only funding projects that set out to develop technologies that complement, rather than substitute, for human workers.

And so, the idea of directed technological progress presents a new set of possibilities for policy-makers. Alongside the traditional preoccupation with increasing the amount of technological progress that takes place—how to get *more* of it—it provides a formal framework for thinking about influencing the quality of technological progress—how to change its *direction*—as well. In this section, I have explored its practical consequences from the point of view of automation and the labour market. But it is important to note that the basic insight—that it is possible to shape the incentives that people face and, in turn, change the direction of technological progress—applies in other domains as well. The opportunity is not simply to direct technological progress with respect to the threat of automation, or indeed the climate as before, but with respect to other things that we might value—to pursue, for instance, new technologies that support our political institutions rather than degrade them, or nurture local places and communities rather than undermine them (the prospect of redirecting technological progress in other domains is the focus of Susskind (2024)).

IV. The limitations of policy

The idea of using policy to influence the way in which technical change is directed has great promise. However, it also faces substantial limitations. These tend to be underexplored by the idea's promoters. And in this section, I want to capture two of the most important types of limitations: the technical and the political.

The technical limits to directed technical change reflect the fact that, while in theory it is plausible to believe that one can clearly distinguish between the different causes of technological progress, in practice those sorts of demarcations are extremely hard to make. And to see this in action consider, again, the impact of technological progress on the labour market. Here, the starting point for those who want to redirect technological progress is the tight conceptual distinction between those technologies that substitute for human workers, reducing the demand for their efforts, and those that complement human workers, increasing the demand for their efforts at unautomated tasks. However, there are both *ex ante* and *ex post* difficulties in making this distinction.

The *ex ante* technical difficulty is clear: it is hard to predict whether a technology will help or harm workers. Indeed, even those technologies that, before adoption, might have appeared to have a straightforward effect in either direction, have not turned out that way. Take, for example, the automated teller machine, or ATM. This is a technology that was intended to displace bank tellers from the task of handing over cash (this case study is set out in Susskind (2020*a*)). They became very common, over a short period of time: in the US, for instance, the number of ATMs more than quadrupled from the late 1980s to 2010. One might have reasonably expected *ex ante* a substantial fall in the number of tellers working in American banks. Yet in practice, the number of tellers rose during the same period in the US, by as much as 20%.

Why was the *ex ante* expectation that ATMs would substitute for human workers mistaken? Because ATMs also complemented workers, albeit in a spread of harder to discern ways: freeing up tellers to focus on other services like face-to-face support and financial guidance, which may have boosted demand through better service; reducing the cost of running branches, allowing banks to draw more demand through lower prices; contributing in part to the productivity gains that not only increased US incomes, leading to more demand for financial services and the remaining tellers (Susskind, 2020*a*). All of these channels are plausible, but they are hard to disentangle. (This case also points to another technical difficulty: that a single technology can have both helpful and harmful effects at the same time, since some are direct effects and the others are more indirect.)

The *ex post* technical difficulty with technical change which is directed in a particular way is a subtler one: the impact of a new technology—in this case, whether it helps or harms workers—can change over time. Consider another technology, the sat-nav systems (Susskind, 2020*a*). In the beginning, these appeared to complement human workers: helping drivers to get from A to B more efficiently and to navigate on entirely unfamiliar roads. But this helpful property will only remain the case so long as human beings remain in the driving seat. In a world of driverless cars, these technologies will have reinforced the effect of software that has displaced human drivers, hurting them instead.

The history of chess playing provides another illustration of the *ex post* technical difficulty. In 1997, the world chess champion, Garry Kasparov, was beaten by a computer system owned by IBM, known as Deep Blue. But this famous victory did not herald the end of human chess playing, substituting for human beings at the task as some of the more pessimistic might have imagined. Instead, it turned out that a good human player, working in a team with this sort of chess-playing machine, could beat any chess-playing machine operating alone: in short, this new technology appeared to complement human beings, making them even better at the task of playing chess. Kasparov approvingly called this 'centaur chess'. (And from this case, he drew an ill-fated conclusion for thinking about the impact of technological progress across the entire economy.) Today, though, Kasparov's centaur is no more. Chess-playing machines, such as AlphaZero, are so capable that human players add nothing to any potential partnership (Susskind, 2020*a*). In short, what began as a technology that complemented good human players now substitutes for them.

Alongside these technical limits, are the political ones: that even if it were possible to make *ex ante* distinctions between the different effects that a new technology might have, and even if there were *ex post* stability with respect to those effects, it might nevertheless still not be feasible to implement them in practice. The challenge of climate change captures this limitation. Here, technological progress, as noted before, can exhibit a bias towards either 'dirty' or 'clean' technologies: the promise of directed technical change is that policy-makers can change the incentives that individuals face and, as a result, encourage the development of the latter not the former. However, the limitation with this broad idea is far less to do with the technical limits, and far more to do with political ones.

For some time, policy-makers have had a spread of possible interventions at their disposal which could in theory reduce emissions sufficiently to meet basic climate ambitions: for instance, a carbon tax that properly captures the damage done by emissions—estimated at around $100 per metric tonne by 2030—would be enough to avoid catastrophic climate change. Put in terms of the idea of directed technological change, a strong enough incentive would help to redirect technological progress, away from dirty technologies to clean ones. Yet in practice, political leaders have not adequately adopted it, or other complementary interventions. And the reason they have failed is not because it is hard to distinguish between clean and dirty technologies or, in turn, that this distinction is unstable over time, but for political reasons: a combination of insufficient domestic political will—those individuals damaged by this transition have been too powerful, and inadequate international cooperation—those countries damaged by this transition have also been too powerful.

The result is that, in most countries, the strategy with respect to addressing the challenge of climate change is some combination of *mitigation*, using the tools at our disposal to try and redirect technological progress, but also *adaption*, recognizing that the political constraints on action limit our collective capacity to respond to the challenge in practice as theory demands. Though this distinction—between mitigation and adaption—is often appealed to in the setting of climate change, it is nevertheless a useful way of thinking about the consequence of the fact that our capacity to deal with the challenges presented by technological progress is limited as well.

V. Conclusion: to what goal?

In this paper, I have explored the idea of directed technological change: its history, its modern form, its practical relevance and usefulness, and its technical and political limitations. However, there is one significant issue that I have not explored: the possibility of redirecting technological progress begs a deeper question—in what direction *should* we heading? 'But in contemplating any progressive movement,' wrote the philosopher John Stuart Mill, 'the mind is not satisfied with merely tracing the laws of the movement; it cannot but ask the further question, to what goal?' (Mill, 1848). In closing, I want to explain why this issue—about ultimate ends, or goals—is so important.

In certain cases, there may be agreement about what those ultimate ends ought to be: avoiding catastrophic climate change, for instance, is likely to be an end that most people would agree upon as a sensible shared goal. And so, with that in mind, it follows that we should do what we can to redirect technological progress in that direction. But for other challenges presented by technological progress, the ultimate end that we ought to be seeking is far less clear.

Consider, for instance, the impact that technological progress has on the labour market. When leading economists explore the challenge of 'excessive automation', there is often an assumption—at times explicit, but often implicit—that we ought to avoid a world with less work for people to do (Susskind, 2023). This view—that there exists a healthy relationship between work and meaning—contrasts with the traditional view of work in the formal economic literature, where the textbook model of individual labour supply instead assumes that work is necessarily a source of disutility, and that effort is only provided in return for a compensatory wage. But in public commentary, the opposite view is often articulated.

Why? Because work is not simply considered by them to be a source of income—which could be provided to those without work through a new redistributive mechanism—but because it is also taken to be a unique source of meaning and purpose: '[m]ost economists are concerned about how we allocate jobs and underneath that concern lies a belief that work matters independently of the earnings that are generated by the work' (Stevenson, 2019).

One can see this belief more clearly in the public commentary provided by leading economists: David Autor, for instance, has stated that '[i]dleness is a terrible thing' whereas '[w]ork gives people's lives structure and meaning'; Daron Acemoglu, has argued that 'it is good jobs, not redistribution, that provides people with purpose and meaning in life (Autor in Wellisz, 2017; Acemoglu, 2019).

For those who agree with these economists—that work is a unique source of meaning and purpose—it follows that a world with less work is not an ultimate end which we ought to welcome: we might be able to provide displaced workers with an income through a new redistributive mechanism, the argument goes, but we can never replace that lost sense of meaning and purpose. If that is right, then it follows that we ought to try to redirect technological progress in such a way that individuals are encouraged to develop new technologies that complement, rather than substitute for, human workers, redirecting technological progress towards a world with sufficient work.

However, there are likely to be many others who may not share the same beliefs as these economists. Suppose you are suspicious of the conflation of work and meaning: you might take strength, for instance, from polling which suggest that most people do not gain a sense of meaning or purpose from their paid work (in the US, for instance, only 50% say they get a sense of identity from their job); or, alternatively, that many people can and do seek meaning outside of paid work (in the UK, for instance, 15 million people actively volunteer, about half the total number of paid employees; Susskind (2021)). And with that in mind, you may find the prospect of a world with less work far more attractive: we could not only provide an income to those without work through redistribution, but also provide richer opportunities to find meaning and purpose through other valuable activities.

The point here is not to take a position on the relative merits of these positions with respect to the desirability of a world with less work. Instead, it is to show that reasonable people might disagree on whether such a world is an end that we ought to pursue, and direct technological progress towards, or to avoid, and turn technological progress away from. Most importantly, the same indeterminacy about which ends we ought to value is likely to characterize many other challenges presented to us by technological progress. Deciding what ends we collectively ought to value, and how we ought to weigh up competing ends, introduces some of the hardest questions that we can ask one another. They are likely to invite immense disagreement. But they are, in the final analysis, exactly the sorts of questions with which we must grapple if we are to take the idea of using policy to direct technological change seriously—as indeed we should.

References

Acemoglu, D. (1998), 'Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality', *The Quarterly Journal of Economics*, 113(40), 1055–89. https://doi.org/10.1162/003355398555838.

— (2002), 'Directed Technical Change', *Review of Economic Studies*, 69, 781–809. https://doi.org/10.1111/1467-937X.00226.

— (2003), 'Factor Prices and Technical Change: From Induced Innovations to Recent Debates'in Aghion, P., Frydman, R., Stiglitz, J. E., and Woodford, M. (eds), *Knowledge, Information, and Expectations in Modern Macroeconomics*, Princeton, NJ: Princeton University Press. https://doi.org/10.1515/9780691223933.

— (2007), 'Equilibrium Bias of Technology', *Econometrica*, 75(5), 1371–409. https://doi.org/10.1111/j.1468-0262.2007.00797.x.

— (2010), 'When Does Labor Scarcity Encourage Innovation?', *Journal of Political Economy*, 118(6), 1037–78. https://doi.org/10.1086/658160.

— (2019), 'Where Do Good Jobs Come From?', *Project Syndicate*, 26 April, https://www.project-syndicate.org/commentary/automation-vs-job-creation-by-daronacemoglu-2019-04.

— (2021), 'Could We and Should We Reverse (Excessive) Automation?'in Blanchard, O., and Rodrik, D. (eds), *Combating Inequality*, Cambridge, MA: MIT Press. https://doi.org/10.7551/mitpress/13469.001.0001.

— Johnson, S. (2023*a*), *Power and Progress*, London: Basic Books.

— — (2023*b*), 'Rebalancing AI', *IMF Finance and Development Magazine*, December.

— — (2019), 'Automation and New Tasks: How Technology Displaces and Reinstates Labour', *Journal of Economic Perspectives*, 128(6).

— — (2020), 'The Wrong Kind of AI? Artificial Intelligence and the Future of Labour Demand', *Cambridge Journal of Regions, Economy and Society*, Cambridge Political Economy Society, 13(1), 25–35. https://doi.org/10.1093/cjres/rsz022.

— Manera, A., and Restrepo, P. (2020), 'Does the US Tax Code Favor Automation', *Brookings Papers on Economic Activity*, 231–300, Spring 2020. https://doi.org/10.1353/eca.2020.0003.

— Aghion, P., Bursztyn, L., and Hemous, D. (2011), 'The Environment and Directed Technical Change', GRASP Working Paper, 21 October.

— (2012), 'The Environment and Directed Technical Change', *American Economic Review*, 102(1), 131–66. https://doi.org/10.1257/aer.102.1.131.

Aghion, P., and Howitt, P. (1992), 'A Model of Growth through Creative Destruction', *Econometrica*, 60(2), 323–51. https://doi.org/10.2307/2951599.

Allen, R. (2009*a*), 'The High-wage Economy of Pre-industrial Britain'in *The British Industrial Revolution in Global Perspective*, Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511816680.

— (2009*b*), 'Why Was the Industrial Revolution British?' *VoxEU* Column, 15 May. https://doi.org/10.1017/CBO9780511816680.

— (2011), 'Why the Industrial Revolution Was British: Commerce, Induced Invention, and the Scientific Revolution', *The Economic History Review*, 64(2), 357–84. https://doi.org/10.1111/j.1468-0289.2010.00532.x.

Autor, D. (2015), 'Why Are There Still So Many Jobs? The History and Future of Workplace Automation', *Journal of Economic Perspectives*, 29(3), 3–30. https://doi.org/10.1257/jep.29.3.3.

— Katz, L., and Krueger, A. (1998), 'Computing Inequality: Have Computers Changed the Labor Market?', *The Quarterly Journal of Economics*, 113(4), 1169–213. https://doi.org/10.1162/003355398555874.

— Levy, F., and Murnane, R. (2003), 'The Skill Content of Recent Technological Change: An Empirical Exploration', *The Quarterly Journal of Economics*, 118(4), 1279–333. https://doi.org/10.1162/003355303322552801.

Bekman, E., Bound, J., and Machin, S. (1998), 'Implications of Skill-biased Technological Change: International Evidence', *The Quarterly Journal of Economics*, 113(4), 1245–79.

Bloom, G. (1946), 'A Note of Hicks's Theory of Invention', *American Economic Review*, March, 83–96.

Bound, J., and Johnson, G. (1992), 'Changes in the Structure of Wages in the 1980s: An Evaluation of Alternative Explanations', *The American Economic Review*, 82(3), 371–92.

Card, D., and Lemieux, T. (2001), 'Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-based Analysis', *The Quarterly Journal of Economics*, 116(2), 705–46. https://doi.org/10.1162/00335530151144140.

Clarke, R., and Summers, L. (1980), 'The Labour Scarcity Controversy Reconsidered', *The Economic Journal, Royal Economic Society*, 90(357).

Dorling, D. (2021), Slowdown, New Haven, CT: Yale University Press.

Drandakis, E., and Phelps, E. (1966), 'A Model of Induced Invention, Growth and Distribution', *The Economic Journal*, 76, 823–40. https://doi.org/10.2307/2229086.

Goldin, C., and Katz, L. (1998), 'The Origins of Technology-Skill Complementarity', *The Quarterly Journal of Economics*, 113(3), 693–732. https://doi.org/10.1162/003355398555720.

Goldin, C., and Katz, L. (2008), *The Race Between Education and Technology*, Cambridge, MA: Harvard University Press.

Grossman, G., and Helpman, E. (1991), 'Quality Ladders in the Theory of Growth', *The Review of Economic Studies*, 58, 43–61. https://doi.org/10.2307/2298044.

Habakkuk, H. J. (1962), *American and British Technology in the Nineteenth Century*, Cambridge: Cambridge University Press.

Harrod, R. (1948), *Towards a Dynamic Economics*, London: Macmillan.

Hicks, J. (1932), *The Theory of Wages*, London: Macmillan.

Kennedy, C. (1964), 'Induced Bias in Innovation and the Theory of Distribution', *The Economic Journal*, 74(295), 541. https://doi.org/10.2307/2228295.

Lucas, R. (1988), 'On the Mechanics of Economic Development', *Journal of Monetary Economics*, 22(1), 3–42. https://doi.org/10.1016/0304-3932(88)90168-7.

Maddison, A. (2001), *The World Economy: A Millennial Perspective*, OECD.

Mill, J. S. (1848), *Principles of Political Economy*, London: John W. Parker.

Rothbarth, E. (1946), 'Causes of the Superior Efficiency of USA Industry as Compared with British Industry', *The Economic Journal*, 56(September), 383. https://doi.org/10.2307/2226046.

Romer, P. (1986), 'Increasing Returns and Long-run Growth', *Journal of Political Economy*, 94(5), 1002–37. https://doi.org/10.1086/261420.

— (1990), 'Endogenous Technological Change', *Journal of Political Economy*, 98(5), S71–S102. https://doi.org/10.1086/261725.

— (1992), 'Two Strategies for Economic Development: Using Ideas and Producing Ideas', *Proceedings of the World Bank Annual Conference on Development Economics*, 63–91.

— (1993), 'Idea Gaps and Object Gaps in Economic Development', *Journal of Monetary Economics*, 32(3), 543–73. https://doi.org/10.1016/0304-3932(93)90029-F.

Salter, W. (1966), *Productivity and Technical Change*, Cambridge: Cambridge University Press.

Samuelson, P. (1965), 'A Theory of Induced Innovations Along Kennedy–Weisacker Lines', *The Review of Economics and Statistics*, 47, 444–64. https://doi.org/10.2307/1927763.

Solow, R. (1956), 'A Contribution to the Theory of Economic Growth', *The Quarterly Journal of Economics*, 70(1), 65–94. https://doi.org/10.2307/1884513.

Solow, R. M. (1957), 'Technical Change and the Aggregate Production Function', *Review of Economics and Statistics*, 39(3), 312–20.

Stevenson, B. (2019), 'Artificial Intelligence, Income, Employment, and Meaning', ch. 7in A. Agrawal, J. Gans, and A. Goldfarb (eds), *The Economics of Artificial Intelligence: An Agenda*, Chicago, IL: Chicago University Press.

Susskind, D. (2016), 'Technology and Employment: Tasks, Capabilities, and Tastes', DPhil Thesis, Oxford University.

— (2019), 'Re-thinking the Capabilities of Technology in Economics', *Economics Bulletin*, 39(1), 280–88.

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

— (2020*b*), 'Work in the Digital Economy'in Skidelsky, R., and Craig, N. (eds), *Work in the Future*, London: Palgrave Macmillan. https://doi.org/10.1007/978-3-030-21134-9.

— (2020*c*), 'A Model of Task Encroachment in the Labour Market', Oxford University Working Paper.

— (2021), 'A World With Less Work', *Boston Review*, May.

— (2022), 'Technological Unemployment'in Bullock, J., *et al.* (eds), *The Oxford Handbook of AI Governance*, Oxford: Oxford University Press. https://doi.org/10.1093/oxfordhb/9780197579329.001.0001.

— (2023), 'Work and Meaning in the Age of AI', Brookings Institution Working Paper, January.

— (2024), *Growth: A Reckoning*, London: Allen Lane. https://doi.org/10.4159/9780674297050.

Swan, T. (1956), 'Economic Growth and Capital Accumulation', *Economic Record*, 32(2), 334–61. https://doi.org/10.1111/j.1475-4932.1956.tb00434.x.

Swan, T. (1960), 'Economic Control in a Dependent Economy', *Economic Record*, 36(73), 51–66.

Temin, P. (1966), 'Labor Scarcity and the Problem of American Industrial Efficiency in the 1850s', *The Journal of Economic History*, 26(September), 277–98. https://doi.org/10.1017/S0022050700068388.

Wellisz, C. (2017), 'Late Bloomer', *Finance and Development*, 54(4), Washington, DC, International Monetary Fund.
