

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/146765>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Automation, Unemployment, and Taxation

Abstract: Automation can bring the risk of technological unemployment, as employees are replaced by machines that can carry out the same or similar work at a fraction of the cost. Some believe that the appropriate response is to tax automation. In this paper, I explore the justifiability of view, maintaining that we can embrace automation so long as we compensate those employees whose livelihoods are destroyed by this process by creating new opportunities for employment. My contribution in this paper is important not only because I develop a theoretical framework that we can use to resolve this urgent policy dispute – a dispute that have been discussed extensively by labour economists and policymakers, but largely neglected by political philosophers – but also because my analysis sheds lights on a wider range of controversies relating to the moral and political importance of unemployment.

Keywords: unemployment, technological change, labor economics, social justice, taxation

1. Introduction

Automation occurs when a new technology or a new use of an existing technology eliminates the need for labour in at least one part of the production process. Changes of this kind can bring the risk of technological unemployment, as employees are replaced by machines that can carry out the same or similar work at a fraction of the cost. We might expect the rate of this trend to quicken as new technologies become available, their cost decreases, and the cost of labour rises as a result of minimum

wage legislation and health and safety laws (Estlund 2018). Because of this, the corresponding risk of technological unemployment may be a considerable one. In the words of Erik Brynjolfsson and Andrew McAfee,

Rapid and accelerating digitization is likely to bring economic... disruption, stemming from the fact that as computers get more powerful, companies have less need for some kinds of workers. Technological progress is going to leave behind some people, perhaps even a lot of people, as it races ahead (Brynjolfsson and McAfee 2014: 11).

Among experts, there's disagreement about the long-term effects of automation. Some foresee a sizeable risk of mass unemployment, sustained by the fact that people will become unemployed at a pace quicker than we can find new uses for their labour. Though it's these sensational predictions that grab the headlines – bolstered by celebrity endorsement from figures such as Bill Gates, Elon Musk, and Stephen Hawking – this remains a minority view among labour economists.¹ Even if it's a mistake to rule out the possibility of mass unemployment altogether, there are at least two reasons to think that it's no more than a tail risk.

First, the prediction of mass unemployment involves a failure to appreciate the ways in which some new technologies complement labour, rather than substitute it (Acemoglu and Restrepo 2018; Autor 2015). In all likelihood, medical technology

¹ The study that is most commonly cited in support of the likelihood of mass unemployment is Frey and Osborne (2017). However, as far as I can tell, little in this study supports of the prediction of mass unemployment. For a nice overview of the state of the literature, see Autor, Mindell, and Reynolds (2020).

will not replace doctors; it will merely assist their decision-making. In some cases, the greater efficiency that results from automation may even increase the demand of labour. For example, as X-ray generators become more affordable, hospitals purchase more of them and so the demand for radiologists rises.

Second, the prediction of mass unemployment involves a failure of imagination regarding the new forms that employment will take. Since 1900, the proportion of the United States' workforce employed in agriculture has dropped from 41% to 2% largely because of automated machinery. But crucially, many of the kinds of individuals who would have worked in that industry a century ago are now employed in industries that didn't exist back then. Indeed, more than 60% of the occupations in the United States today had not been 'invented' by 1940, and a substantial portion of these may have been barely imaginable, even to labour market experts (Autor, Mindell, and Reynolds 2020: 10-11). The same may be true for the future. The fact that we're unable to foresee the new forms that employment will take doesn't justify the conclusion that there won't be new forms of employment to replace those with which automation dispenses (Autor 2015).

Despite this, automation still poses some risk of unemployment in the short- and medium-term as the labour market adjusts to the availability of new and cheaper technologies. This unemployment results from a mismatch between the new skills that employers desire (say, programming skills or personal service skills) and the increasingly outdated skills that most individuals possess (say, manufacturing skills). One effect of this is a growth in redundancies, as technology is introduced that is too costly for existing employees to learn to operate. Another consequence is that growing firms will tend to invest in more machinery rather than to expand their

labour force. It's little consolation to the cashier whose role is automated to be told that predictions of mass unemployment are unfounded, and that the economy will adjust in the long-run.²

We might expect the worst hit by this trend to be those with the highest wages, since it is by replacing these employees that firms stand to make the greatest savings. All else being equal, a firm gains much more by replacing its creative director on a high salary than it does by replacing the lift operator on a pittance. However, this reasoning overlooks two sets of facts. First, the tasks performed by low- and medium-earning workers are often more straightforwardly automated than that of high-earning workers. It's much cheaper to automate the task performed by the lift operator than it is to do so for those tasks performed by the creative director. Because of this, those with the highest wages tend to be the most immune to displacement at the hands of automation (Acemoglu 2002; Acemoglu and Autor 2011).

Second, we must be careful about what we mean by 'the worst hit'. Even if automation were to lead to high levels of redundancies among the highest earners, it makes a moral difference that, in virtue of their skills, these individuals tend to remain well-equipped to find new lines of work. Coupled with the likelihood of generous savings and ownership of their homes, very few of the highest earners would face any serious risk of poverty even if they were to become unemployed. By contrast, the reverse is true for the lowest earners: even if their risk of unemployment as a result of automation is low, the threat is much more serious given that these individuals tend to be less employable and that even temporary periods of unemployment can be

² This is in the spirit of John Maynard Keynes's claim that, 'in the long-run, we're all dead' (1923: 80).

devastating. This result is significant because, if our economic system is to be just, it is morally imperative that it serves the interests of our society's least advantaged members.

One response to these anxieties is to introduce policies to decelerate the pace of automation and, therefore, slow down the rate of which employees are replaced by machines. In pursuing this strategy, we exploit the fact that the trend towards automation isn't a natural or unstoppable process. It's a product of our political decision-making regarding the design of our economic institutions.

This move comes in as many varieties as there are ways in which to vary the rate of automation. One influential idea is to introduce *automation taxes* (sometimes also called *robot taxes*), which increase the cost to firms of investing in the equipment that enables automation (McGoogan 2017; Reuters 2017). South Korea has already announced plans of this kind, and the policy has been the subject of heated political debate elsewhere, including in the European Parliament.³ One of the chief draws of this policy is that it financially penalizes automation, and thereby reduces the associated risk of unemployment. We can call this the *disincentive argument* for automation taxes.

In addition to this, automation taxes also serve a second function, namely as a mechanism through which to raise tax revenue. More precisely, these taxes offer a way in which to swell the pool of funds available to governments to soften the blow of automation on those who lose out as a result of technological change. We can call this the *revenue argument* for automation taxes. This argument operates independently

³ For comprehensive discussion of various proposals of this kind, see Bogenschneider (2020).

of the disincentive argument, such that neither's success as a justification for automation taxes depends on the other. It's for this reason that we must evaluate these two arguments separately.

My main aim in this paper is to examine the justifiability of automation taxes, focusing mainly on the disincentive argument and the revenue argument. My finding is that, subject to some qualifications, neither argument provides a sound basis for defending automation taxes. Near the end of the paper, I then briefly consider a technical argument – the *tax neutrality argument* – which sees automation taxes as correcting a defect with the prevailing tax regimes in many states. I conclude that we should be sceptical of this argument too and that, even if successful, its implications are more rather limited than those associated the disincentive argument and the revenue argument.

My contribution in this paper is important not only because I develop a theoretical framework that we can use to resolve a series of urgent policy disputes – disputes that have been discussed extensively by labour economists, tax lawyers, and policymakers, but largely neglected by political philosophers – but also because my analysis sheds lights on further controversies relating to the moral and political importance of unemployment.

2. Taxing Automation

As a preliminary to my discussion, let's consider a simple argument against the use of automation taxes that I consider a non-starter. It appeals to the idea that the benefits of automation are so morally valuable that they count decisively against all attempts to justify policies that decelerate the pace of this process. One way in which

to support this verdict is to invoke the *utilitarian criterion*, according to which a given policy is just when and because it does better than all rival policies at maximizing aggregate benefits and minimizing aggregate burdens across the entire population of the relevant political community. On this view, though anti-automation policies may have some appeal, specifically to those employees who're most at risk of being replaced by machinery, these reasons are defeated by our reasons to promote productive efficiency.

One familiar way to express this idea is by using the language of Kaldor-Hicks efficiency (Hicks 1939; Kaldor 1939). A policy is efficient in this sense when there's no way to improve one person's circumstances such that her gains are sufficient to compensate those who lose out from her improvement. What matters for Kaldor-Hicks efficiency is that those who benefit from a change in policy *could* use their winnings to compensate those who lose; it's not necessary for them actually to do so. Thus, policies that promote automation are Kaldor-Hicks efficient if the gains to those whose situation is improved are sufficiently large that they could compensate those whose circumstances are worsened by automation, though they need not actually do so.

One problem with the utilitarian criterion is that it's insensitive to the distributive effects of our decisions. In particular, providing that the benefits and burdens are of the same size in both cases, it treats policies that burden the least advantaged in order to assist the most advantaged identically to those policies that do the opposite. According to the utilitarian criterion, all that matters is whether the policy maximizes aggregate benefits and minimizes aggregate burdens. Thus, whether those who're affected are destitute or billionaires makes no difference to the justifiability of a

policy, except in so far as these happen to influence the size of the benefits and burdens in question (say, because of income's diminishing marginal returns).

But, this feature of the view is powerfully counter-intuitive. It must make a moral difference that the worst hit by the threat of unemployment are those with lower wages, and that most of the rewards will be reaped by members of more advantaged groups, such as higher earners, business owners, and shareholders. This is important because, as John Rawls famously notes, we may not justify policies "on the grounds that the hardships of some are offset by a greater good in the aggregate" (1999: 13). Doing so may be expedient, but it cannot be just.

If correct, the upshot of these (very) brief remarks is that we should resist the conclusion that the benefits of automation are so morally valuable that they count decisively against all attempts to justify policies that decelerate the pace of automation. The fact that policies that promote automation might satisfy the utilitarian criterion doesn't guarantee their justifiability. Instead, this is a complex matter that depends, among many other things, on how rival policies affect the prospects of the least advantaged. With this in mind, we now turn to the first argument in favour of automation taxes.

3. The Disincentive Argument

According to proponents of the disincentive argument, we should introduce automation taxes to discourage firms from choosing technology over labour. In doing so, this policy protects us from the threat of unemployment that would otherwise occur.

The strength of these reasons to oppose automation depends largely on the extent to which we can effectively compensate those whose employment is displaced through this process, so as to reduce the disadvantage that these individuals would otherwise suffer. If compensation were very costly, then these individuals' complaints against automation would be strong ones that are difficult to overcome. In this case, the disincentive argument is especially appealing, since it responds to the threat of unemployment by averting it rather than by compensating its victims. Borrowing a distinction from climate justice, we might say that proponents of the disincentive argument respond to the threat of unemployment by *mitigating* that threat rather than *adapting* to it. By contrast, if compensation were less costly, then the force of these complaints would dissipate – or, more precisely, it would dissipate so long as compensation is actually on offer.

Before proceeding, it may be useful to clarify how I use the term 'compensation'. Though it's perhaps natural to understand it in a purely pecuniary way – that is, as financial compensation – I adopt a broader definition that includes non-pecuniary elements, such that we might compensate an individual by providing her with benefits in kind, perhaps including new employment opportunities (as I discuss below).

Moreover, though recipients of compensation are sometimes victims of *wrongdoing*, this isn't always the case. Indeed, in some cases, the guarantee of compensation can even justify conduct that would otherwise be wrongful. For example, it can sometimes be permissible for an individual to cause noise pollution on the condition that she compensates her neighbours as a result, where it would be wrongful to cause such noise pollution in the absence of compensation. In this case, compensation isn't

a response to wrongdoing; its guarantee can render the conduct morally justifiable (Nozick 1974: 174-82). This point is important since it reveals that there's nothing inherent in the idea of compensation, as I use the term, that makes it a less-than-perfect solution, in the sense that its need arises only when individuals fail to respect others' rights or otherwise act wrongfully. As we shall see, these ideas inform my critique of the disincentive argument.

One way in which we can compensate the victims of automation is by providing them with unemployment benefits, whereby those who're unemployed can claim regular cash payments from the state to assist them in difficult times. Ordinarily, claimants' entitlements to unemployment benefits are conditional on them exhibiting a willingness to work. Therefore, these payments are normally denied to individuals who're not actively seeking employment, those who decline job offers, and those who have recently resigned. These schemes may vary in a number of dimensions, including in the level of the payments and the conditions that attach to these, as well as whether they persist indefinitely or cease after a specified number of years.⁴

As a mechanism through which to compensate those who lose out as a result of automation, unemployment benefits are attractive because regular cash payments increase the income available to individuals to pursue their goals. It's true that an individual with no disposable income may enjoy the formal freedom to go to the

⁴ Some insist that a policy of basic income is preferable to that of unemployment benefits. But the strongest case for this depends on arguments that are incidental to our purposes, such as the fact that basic income improves the bargaining power of the least advantaged. Because of this, it makes sense to set this issue aside, though I discuss the appeal of basic income in other work. See Parr (unpublished manuscript).

shops, but that's of little value to her if she can't afford to purchase anything once there. What's appealing about unemployment benefits is that, by providing someone who's unemployed with disposable income, it converts some of her formal freedoms into effective freedoms.

However, it's a mistake to think that unemployment benefits alone can be a satisfactory method of compensation. This is because being employed can be valuable for a host of reasons other than the income that it provides – some of which are fairly obvious, such as that it provides a prime opportunity to develop and exercise our talents, and others of which are less obvious, such as that it tends to improve our physical health. Rather than survey these reasons, I want to focus on one reason of special significance, namely that employment affects *self-respect*, defined in terms of having a secure sense of the worth of one's goals and the confidence to pursue them (Rawls 1999: 386). There are two aspects to this, each relating to a fact about our moral psychology.

First, an individual's self-respect depends on her status in the eyes of others as being able to live as a normal functioning agent (Scanlon 2018: ch. 2; Ci 2014). In a society with mass unemployment, self-respect would not depend nearly as much as it currently does on being employed, since this wouldn't be necessary to being seen by others as a normal functioning agent. But, so long as most people remain in employment, there's reason to think that employment will continue to play this special role in our lives. In this way, the value of being employed increases as the

number of others in employment also increases.⁵ We can call this the *status aspect* of self-respect.

In support of this idea, let's consider the fact that divorce has been a source of shame and stigma that has damaged the self-respect of many people around the world, especially women. One explanation for this is that, where this is the case, a woman who divorces isn't seen by others as able to live as a normal functioning agent. She's seen as atypical in a crucial respect. Of course, this isn't a justification for the shame or stigma; it's an explanation of the phenomenon that relies on a claim about moral psychology. Importantly, though, in some parts of the world, divorce is no longer a source of shame where it once was. In some cases, part of what explains this trend is rising rates of divorce within those communities, such that divorce is increasingly seen as part and parcel of life. This accounts for why divorce tends to be less of a source of shame in communities where rates of divorce are high than in communities where rates are low. This provides some evidence in support of the status aspect of self-respect.

Second, to some degree, an individual's self-respect depends on her having her efforts appreciated by others – that is, it depends on the receipt of external validation. This is because, as Rawls notes, “unless our endeavors are appreciated by our associates it is impossible for us to maintain the conviction that they are worthwhile” (1999: 388). We can call this the *validation aspect* of self-respect. Though it's possible to gain this kind of validation outside of employment, such as in our roles as family members, the signals that convey it may be less reliable in this sphere than in the

⁵ I defend this more fully in Parr (unpublished manuscript).

market. This is because we've special reasons to get along as families, and because talk is cheap. By contrast, in the workplace, we lack these special reasons to get along (or, at least, they are less weighty) and employers bear costs in showing their appreciation of our efforts by paying wages.⁶ This is why the external validation that we attain through employment has particular significance to our self-respect (Jütten 2017: 267-9).

Three features of this analysis are worth highlighting. First, employment can serve individuals' self-respect only if it surpasses some threshold of decency (Shelby 2016: ch. 6). This isn't the case if the work is highly degrading and poorly paid with few prospects of promotion, for example. How to draw the line between decent and indecent employment is a tough issue, which I can set aside here. Second, it's consistent with these claims that there are some individuals whose self-respect does not depend on employment, at least not to any meaningful degree. This is the case for those whose class privilege confers a robust sense of confidence that's not threatened by involuntary unemployment. Third, we can leave open the possibility that the intimate link between employment and self-respect is a regrettable one. This is especially so in practice once we recognise the way in which pro-employment norms are constructed and applied in racist and sexist ways, celebrating the achievements of some demographics and overlooking those of members of other groups (Lamont 2000).

⁶ As further evidence of this, we can consult the testimonies of those who return to employment following an extended period of parental leave, many of whom emphasize a renewed sense of having their efforts appreciated by others. See also Crosby (1987).

Crucially, though, it doesn't follow from these remarks that we should oppose high rates of automation, as maintained by defenders of the disincentive argument. It doesn't even follow that unemployment benefits shouldn't be part of the solution. After all, if the reason that this policy is unsatisfactory is that it only partially compensates those displaced by automation – that is, it compensates them for some losses but not for others – then we can remedy this deficiency by supplementing this policy with additional measures that create new opportunities for employment. Typically, this will involve the use of supply-side policies, whose purpose is to equip individuals with the new skills that make them desirable to potential employers. Examples of this include investment in education and effective retraining programs. But, it may also involve the use of demand-side policies, whose purpose is to create new demand for labour. For example, governments might subsidize new hires; invest in labour-intensive industries; expand public sector employment; or enact a job guarantee scheme, whereby the state serves as an employer of last resort (see Blank and Card 2000: chs 7-9; Gottschalk and Freeman 1998). Precisely which mix of policies is most effective is a tricky economic issue that's highly context-sensitive, and so I suspect it doesn't admit of any general answer (Atkinson 2015; Brynjolfsson and McAfee 2014: ch. 13; World Bank 2019: 9-11).

At this stage, an advocate of the disincentive argument might seize on this result, insisting that automation taxes should be among the suite of policies available to governments to maintain demand for labour. For example, there may be cases in which, rather than subsidizing new hires, it's better to raise the cost of the alternative by taxing automation. If that were correct, then my line of reasoning could not rule out the justifiable use of automation taxes, at least in principle.

In the first instance, one problem with this move is that it elides the distinction between (a) policies that *preserve* demand for labour, such as automation taxes, and (b) policies that create *new* demand for labour so as to make up for the shortfall that would otherwise exist because of automation. My suggestion is that, subject to a complication that I'll address shortly, (b) is no worse than (a); and, on the plausible assumption that automation taxes bring additional economic costs due to reduced technological innovation, then (b) is preferable to (a). The upshot is that we've reasons to prefer policies that create new demand for labour, rather than to introduce automation taxes.

My response to the disincentive argument relies on the idea that we can protect individuals against the threat of technological unemployment by creating new opportunities for employment. But critics will point out that we must take seriously the possibility that doing this will sometimes be prohibitively costly, especially if the pace of automation is rapid such that large numbers of people are frequently displaced by this process (Estlund 2018: 280-3). The stalling of countless regional economic regeneration programmes, from the North East in the United Kingdom to the Rust Belt of the United States, provide some evidence that creating large numbers of new employment opportunities is easier said than done.

Because of this, we've sound reasons to ensure that the pace of automation does not exceed a rate at which we can compensate all those who lose their work due to automation with new opportunities for employment. Where this is not the case, such that automation becomes a run-away train that could cause unmanageable destruction, the disincentive argument for automation taxes succeeds.

Though I'll not substantiate this empirical claim here, since it's not essential for my current purposes, my sense is that we're some way short of the point at which this condition is satisfied. No doubt, that there are some areas with persistently high levels of unemployment. But, I'm yet to be convinced that this is because automation is causing 'unmanageable destruction', rather than that this is a product of massive under-investment in regeneration programmes.

4. The Adjustment Objection

Defenders of the disincentive argument might attempt to revivify their position by pointing out that, even if we can significantly mitigate the costs of being displaced as a result of automation by creating new opportunities for employment, these policies will not eliminate all of the relevant costs since it remains burdensome to transition to new employment. We can refer to this as the *adjustment justification* for the disincentive argument.

Three factors support this justification. First, there will be some individuals for whom the prospect of re-training and new forms of employment are of little consolation. This is most clearly the case for older employees, including those made redundant later in life, maybe only shortly before their retirement. The upheaval involved in changing jobs at this point in life is a potentially serious burden that we must not overlook. Second, unemployment leaves scars whereby those who return to work following a period of unemployment typically receive lower wages than they would have if they'd remained employed throughout (Arulampalam 2001). Third, even if new opportunities for employment render individuals no worse off, we must register the fact that this process none the less forcibly shifts people to change their plans and lifestyles (Goodin 1989). In other words, even if these policies were to eliminate any

disadvantage, having to transition to a new form of employment would still interfere with life plans in a way that may call for justification.

The adjustment justification gets particular intuitive force from reflecting on cases in which individuals are made redundant after a long period in the same post, and where they have structured various parts of their lives around that. However, this is an increasingly outdated account of people's experience of the labour market. Other things equal, it's now more common for an individual to switch jobs, even to switch industries, multiple times throughout her working life, including only shortly before retirement (World Bank 2019: vii). Moreover, many of us structure the rest of our lives accordingly – and can reasonably be expected to do so – in order to reduce the costs involved in changing jobs, such as by having friends beyond one's colleagues. Because of this, we can expect the adjustment costs associated with transitioning to new employment to be lower for future cohorts of employees. In turn, this weakens (but does not neutralize) the force of the adjustment justification.

Explaining the justifiability of the remaining adjustment costs is a tricky business. It involves pointing out three things. First, the mechanisms that impose these costs also produce considerable benefits for other people. As noted during the discussion of the utilitarian criterion, the mere fact that pro-automation policies maximize aggregate benefits and minimize aggregate burdens isn't sufficient to establish their justifiability, since we also care about the distribution of benefits and burdens. However, it can be reasonable to expect individuals, including those who're among the worst off, to bear *some* costs when doing so greatly improves many others' prospects.

Second, the mechanisms that impose these costs also produce considerable benefits for the very same individuals, in their capacity as consumers. To be sure, even if it were true that, for each of us, automation's benefits outweighed its burdens, this wouldn't be sufficient to establish the justifiability of pro-automation policies. However, if the benefits of automation greatly outweigh its burdens, and this isn't true for alternative arrangements, then this can count strongly in its favour. This might be the case if an individual is made redundant shortly before retirement but, as a consumer, she reaps the considerable rewards of living in an economy that prizes innovation and that distributes its fruits widely. It's easier to justify the burdens that this process imposes on her in these circumstances than when the benefits of automation that she enjoys are much smaller.

Third, the adjustment costs that we're considering may be stubborn, such that it's often very difficult, and perhaps even impossible, to eliminate them without giving up on the fruits of automation, or indeed technological progress more broadly. This is important because it means that we can't divorce our evaluation of the justifiability of imposing these costs from an account of the benefits of automation and technology. This wouldn't be the case if there were two mechanisms through which to realize the benefits of automation that differed only in terms of the presence or absence of these adjustment costs. In that case, it would be unjust to opt for the more costly policy. But alas, that's not the situation in which we find ourselves. (I explore a related idea in more detail in section 6.)

The upshot of these three considerations is that, even if we can't eliminate all of the costs associated with transitioning to new employment, the remaining adjustment costs will not normally be sufficiently serious to justify measures, such as automation

taxes, whose purpose is to reduce the rate at which employees are displaced because of technological change. Instead, these costs can be justifiable given that they're an unavoidable part of a process that produces much larger gains from which everyone benefits, including those who bear those costs.

5. The Revenue Argument

A defender of automation taxes might re-assert their view by appealing to another of this policy's features, namely that it serves to swell the pool of funds available to governments to compensate the victims of automation by creating new opportunities for employment. On this view, the fact that these taxes reduce the pace of automation is an inessential part of the policy – indeed, we might even regard it as regrettable. What matters is that these taxes provide an effective route through which to ascertain funds in an age where rapid automation might severely reduce labour-related tax revenue such as funds raised through income tax. This is the revenue argument for automation taxes (Abbott and Bogenschneider 2018; Bogenschneider 2020; McCredie, Sadiq, and Chapple 2019).

In order for this argument to succeed, we must explain why this tax burden should fall specifically on firms that automate, rather than on other groups in society or on the general public writ large (Estlund 2018: 317). In this section, I consider two possible justifications for this claim, but show that neither is satisfactory. The upshot is that we should resist the revenue argument.

5.1. Harm-Based Taxation

One straightforward argument for placing the tax burden on firms that automate is that these firms are causally responsible for the plight that the individuals who they

replace with machines would otherwise suffer. On this view, automation taxes raise revenue by serving as fees that correct for the fact that the market price of unemployment-producing technology fails to reflect the harmful externalities of automation on the lives of those displaced by this process. This is an example of *harm-based taxation* (McCredie, Sadiq, and Chapple 2019: 655-6).⁷

To illustrate the appeal of this reasoning, it may help to consider a simple case in which the use of harm-based taxation is more familiar:⁸

Pollution: Evi sets up a factory that releases pollutants into a nearby stream, and so contaminates the local water supply. As a result, the yield of Kapo's crops decreases markedly.

In this case, we should tax Evi's conduct at a rate that reflects the costs imposed on Kapo, and use the revenue raised to compensate Kapo, perhaps by subsidizing the cost of fertiliser or by giving him cash transfers. Justified in this way, the environmental tax operates as a fee that corrects for the fact that the market price of Evi's activities fails to reflect the harmful externalities of this process on Kapo and his livelihood. This is what explains why we're justified in raising the revenue through taxes levied on Evi in particular, rather than on other groups in society or on the general public writ large.

⁷ Some authors regard these taxes as Pigouvian, but I prefer to avoid this label since the Pigouvian justification for taxation refers to the socially optimal level of harmful conduct, rather than the moral importance of protecting individuals against those harms. See Pigou (1960 [1932]).

⁸ Though it's used for a somewhat different purpose, I take this simple example from Abel et al. (2021: ch. 13).

It's likely that the use of this environmental tax will disincentivize Evi's conduct, and so reduce or eliminate the need for compensation. However, to the extent that this is a welcome result, this is for reasons linked to the disincentive argument, and so something that we can set aside for the purposes of analysing the revenue argument. What's significant about this case is that it intuitively illustrates the claim that those whose conduct produces the need for compensation should bear the costs necessary to fund it. If these insights were to generalize to the case we're considering, then we'd have a sound basis on which to support automation taxes.

The chief problem with this argument is that the parallel between Evi and firms that automate is a misleading one. This is because, whereas Evi has a general moral duty not to pollute – albeit a duty that she may be able to circumvent if she guarantees compensation for Kapo – firms have no general moral duty to employ staff (or otherwise to ensure that other firms employ those individuals). For this reason, even though firms that automate may be causally responsible for the plight that the individuals who they replace with machines would otherwise suffer, they aren't morally responsible in the same way that Evi is morally responsible for the costs that Kapo would otherwise bear.

Let's compare two expanding start-ups that are identical in every way, including profit levels, except that, whereas Firm X expands by growing its labour force, Firm Y expands through investment in robotic technology. It's counter-intuitive to suppose that Firm Y is liable to a higher level of taxation than other firms on the grounds that it bears special moral responsibility for the disadvantage suffered by those who'd be employed if it had chosen the alternative method of expansion. After all, if what grounds the alleged additional liability is that it could have acted to benefit

these individuals, then it's unclear why these burdens should fall on this firm rather than be shared with everyone who has this or a similar opportunity, including those who refraining from setting up profitable firms when they've the opportunity to do so. To extend additional liability to these individuals would be an even more obviously unwelcome result.

Our story gets more complex when we introduce Firm Z, which is identical to Firms X and Y in every way, including profit levels, except that it grows by laying off some of its employees in order to invest in robotic technology. Perhaps Firm Z *does* bear some special moral responsibility for the fate of the employees it makes redundant.⁹ This explains why it may be wrong for a firm to dismiss an employee without notice and to fail to offer her a redundancy package. But, it's noteworthy that even very demanding duties of this kind are consistent with the judgment that these firms shouldn't have to bear additional burdens for the specific purpose of generating new employment opportunities for those they make redundant, which is the issue we're considering here. With respect to the revenue argument, what's at stake is who should pay the taxes to fund the creation of new employment opportunities in particular, not redundancy packages or other costs in this ballpark.¹⁰

⁹ For an argument that militates against this claim, see Estlund (2018).

¹⁰ I have discussed these cases as if firms' choices and their effects are publicly verifiable, such that we can identify their winners and losers. In practice, things are likely to be much more complicated, since an individual's prospects may be overdetermined by a plurality of factors. This may create additional difficulties for defenders of harm-based taxation, but I set these aside since my response need not invoke such facts.

We can buttress this reply by pointing out that the parallel between Evi and firms that automate misunderstands the role that firms play within the economy and, in particular, the moral responsibility that firms bear from *competition harms*. This term refers to the harms that occur in losing competitions for scarce goods. This is the case when one firm suffers lower profits as a result of another firm's market success. Relevant to our purposes is the fact that competition harms occur when an individual misses out on an offer of employment or is made redundant because the potential employer instead chooses to automate.

Competition harms are a subclass of unintended harms that are morally permissible only if we've good reason to preserve the competition in question (Clayton and Stemplowska 2016; Dworkin 2011: 285-91). In these cases, the upshot is not only that we're permitted to act in ways that lead to competition harm, but also that we don't owe compensation when the harm arises. This result is plausible if the benefits of having a competitive economy are large enough to outweigh the harms in question very greatly. On this view, the reason that the owner of a firm whose success sets back others' interests bears no special responsibility for their worsened position is that she's merely playing her role within a justifiable system of rules. To impose special responsibilities on her would undermine the competition in a way that jeopardizes its capacity to produce significant benefits. For the very same reason, it's a mistake to impose harm-based taxes on firms that automate as a response to the competition harms that this process generates.

5.2. *Moral Desert and Taxation*

Though rarely discussed by academic scholars, especially in economics and tax law, there is a second justification for the revenue argument that appears more commonly

in public debate. It appeals to the alleged fact that firms that choose not to automate tend to make more of a social contribution than firms that do. This is because they perform the public service of offering opportunities for employment, and because they make more of a contribution to the public coffers through things like payroll taxes than those who employ fewer staff. Firms of all sizes are sometimes at pains to draw our attention to these facts, which they might do by highlighting their track-record of hiring large numbers of employees or by emphasizing how their business stimulates the economy in a way that creates new employment opportunities.¹¹

This is important because some believe that those who make a greater social contribution *deserve* to fare better than those who make a lesser social contribution. And if correct, we might exploit this fact to justify why firms that do not automate should fare better than those that do automate. This argument has the potential to explain why we are justified in taxing the two kinds of firms at different levels, as would be the case if we used automation taxes.

We should dispute this argument on three grounds, each of which is independently forceful. First, it's unclear that firms that choose to automate tend to contribute less to society in virtue of doing so. While it's true that these firms might not directly employ as many staff, the technology that they purchase must come from somewhere. This demand for machinery creates new demand for employment and provides incentives for the kind of innovation from which we all benefit (Autor,

¹¹ Indeed, even Amazon, which is famed for its automation, boasts that its services have helped to create over a million jobs in small- and medium-sized businesses in the United States. See Wilke (2020).

Mindell, and Reynolds 2020: 11). For this reason, it's difficult to make any general claims about the relative social contributions of firms that do and don't automate.

Second, our judgments about desert are most powerful when those contributing to society do so for reasons that aren't self-interested. For example, it's intuitive that someone who saves a child from a burning building out concern for the child's life is more deserving than someone who acts identically but in order to attain celebrity status. This is important because firms are standardly motivated by the self-interested desire to maximize profits, rather than by some altruistic concern for others' employment opportunities. Their claims to privileged treatment on grounds of desert are therefore much weaker than might otherwise be the case.

Third, and most fundamentally, I'm inclined to think that we should oppose the idea that an individual's economic entitlements should depend on her deservingness. The person who selflessly saves a child from a burning building may deserve some things, such as our praise and perhaps some public recognition of her heroism. She may also be entitled to compensation for the harms that she suffered in carrying out the rescue. But, it's far from obvious that she should enjoy greater economic prospects than the rest of us, who did not and who would not enter the building, such that it's unjust for her to fare no better than any other law-abiding individual. Obviously, this reply raises issues that are deeper and more general than the revenue argument (Scanlon 2018: ch. 8). I can't defend this suspicion here, but it bears noting that my views aren't especially eccentric in this regard, at least within contemporary political philosophy.

To summarize, the revenue argument fails because there's no reason why the taxes needed to raise funds for creating new opportunities for employment should fall specifically on firms that automate. Precisely how we should raise these funds is a

matter about which we can remain ecumenical. One option is to raise corporation taxes, which have declined significantly in recent decades. For example, in the United Kingdom, they dropped from 52% in 1982 to 19% in 2017; and in the United States, they dropped from 40% in 1987 to 21% in 2018. Another option is to introduce inheritance or wealth taxes (Piketty 2014: 515-39). The optimal level of these taxes is not an issue that I propose to discuss here. Rather, all that matters for present purposes is that, when raising revenue, we lack a reason to discriminate between firms based on whether or not they automate and, therefore, we lack a justification for the revenue argument.

6. The Tax Neutrality Argument

In this section, I shall consider the tax neutrality argument for automation taxes, which is rather different in character from the arguments that we have seen so far. Its case starts with the fact that fiscal regimes, including those operating in the United States and the United Kingdom, aggressively subsidize firms' investment in capital (via tax credits and accelerated amortization) and aggressively penalize firms' investment in labour (via payroll taxes). This means that, as things stand, these regimes are heavily biased in favour of automation.¹² Proponents of the tax neutrality argument contend that these incentives are distortionary; that this leads to excessive levels of automation (relative to the efficient level that would occur under perfect competition); and that we should employ automation taxes to correct for these problems. The relevant point is not that our tax system should favour investments in labour over investments in capital. Rather, it's that our tax system should be neutral,

¹² Not everyone agrees with this characterization. For example, see Hemel (2020).

or at least much more neutral than at present, with respect to these two factors of production (Abbott and Bogenschneider 2018; Acemoglu, Manera, and Restrepo Working Paper; Bogenschneider 2020; Soled and Thomas 2018).

To highlight the moral concern that animates this argument, let's consider the fact that some technologies are little or no better in terms of productivity than the labour that they replace. Arguably, self-checkout machines are such an example (Acemoglu, Manera, and Restrepo 2020). This is because this technology is expensive to introduce and to maintain, and because the average shopper is probably slower than a skilled cashier at scanning groceries and other such goods. Yet, because of the tax advantages of investing in capital, firms have a financial reason to purchase these machines rather than hire staff to work the tills. The upshot is an economic system that incentivizes excessive automation, such that firms are encouraged to adopt technologies that displace workers but that bring few productivity gains. Viewed with cases such as these in mind, automation taxes are attractive because they reduce firms' incentives to over-spend on capital and to under-spend on labour.

The chief concern with this approach, though, is that taxing automation risks hindering innovation (Mazur 2019). This is morally problematic since the use of some technologies has sizeable positive externalities from which many individuals stand to benefit, especially members of future generations. The promise of these benefits give us weighty reasons to subsidize technological progress (at least so long as the rate of unemployment does not become unmanageable), which in turn implies that we have sound externality-based reasons to abandon tax neutrality. To be sure, advocates of the tax neutrality argument are correct to point out that, in the absence of automation taxes, firms will sometimes have an incentive to adopt technologies

that are little or no better in terms of productivity than the labour that they replace. But this is the price that we must pay to enjoy the fruits of an economy that prizes innovation.

No doubt, some critics will allege that this analysis exaggerates the positive externalities of automation, at least in comparison with those associated with hiring employees. And this concern would be fair if a consequence of a firm's investment in capital were that its employees were displaced into unemployment. But this should not be the case since, as I explained in section 3, my proposal is that we should compensate these individuals with new employment opportunities. This feature of my view is important because the positive externalities of investing in labour are much smaller when this is the fate of those who lose out as a result of automation than when they are condemned to unemployment.

Alternatively, some defenders of automation taxes acknowledge the positive externalities of technological progress but insist that we must tailor these taxes' design such that they apply exclusively, or at least predominantly, to those technologies that are little or no better in terms of productivity than the labour that they replace. This is the view of Daron Acemoglu, Andrea Manera, and Pascual Restrepo, who maintain that these taxes should not "target all forms of capital automation" but instead "focus taxes on automation technologies that are being used for marginal tasks where it does not bring much productivity gain" (Acemoglu, Manera, and Restrepo 2020: 12). I think that this ambition is a sensible one, but I worry that delivering on this commitment requires too optimistic a view of the capabilities of governments and policymakers to predict the full range of effects of various technologies in advance of their adoption, and then to construct a tax regime that is

appropriately sensitive to these considerations. In reality, these tasks are exceptionally difficult ones and probably even impossible most of the time, given the complexity of modern markets. It's telling that Acemoglu, Manera, and Restrepo seem to concede as much when they write, "How would all this work in practice? That's the trillion-dollar tax question" (Acemoglu, Manera, and Restrepo 2020: 13).

What's more, even if supporters of the tax neutrality argument can overcome this challenge, it's significant that their position justifies a much more limited conclusion than that supported by the disincentive argument and the revenue argument. This is because they propose automation taxes as an appropriate means by which to level the playing field between spending on capital and spending on labour. But this line of reasoning no more supports automation taxes than reducing payroll taxes, for example. After all, if the goal is to achieve tax neutrality between these two factors of production, then we can do this *either* by increasing the price of one (via automation taxes) *or* by decreasing the price of the other (via payroll tax cuts). There's nothing in the tax neutrality argument as such that counts in favour of pursuing the first route over the second.

7. Conclusion

What emerges from my discussion of automation taxes is a rival account of how policymakers should respond to the threat of technological unemployment. It has two components. The first is that we should protect individuals against the risk of automation by creating new opportunities for employment. The second is that we should raise the revenue necessary to fund these measures through taxes that don't discriminate between firms on the basis of whether or not they automate.

Critics may be quick to point out that both of these moves are unlikely to happen anytime soon. Governments and the voters who elect and support them have routinely failed to protect individuals against the threat of unemployment, and there's little reason for optimism about the prospects of them doing so in the near future. Moreover, perhaps the most attractive feature of automation taxes is that they provide a politically viable way to combat the effects of lower levels of labour-related tax revenue, as well as unprecedentedly low rates of corporate taxation (Abbott and Bogenschneider 2018).

The point is not merely that my recommended policies are politically unrealistic. The reason these political facts matter is because criticizing the use of automation taxes may be dangerous, given that the conditions necessary for automation to be justifiable are unlikely to be satisfied. Therefore, because of the predicted failure of governments to enact the policies whose use I have defended, perhaps we should favour automation taxes as a next-best solution. After all, given the havoc that automation would wreak if we failed to create new opportunities for employment, automation taxes may be better than nothing (Abbott and Bogenschneider 2018). This line of reasoning is especially attractive when the burdens of the policy fall predominantly on firms that fail to pay their fair share of taxes.

This is a serious concern, and I acknowledge that, as a matter of political strategy, we should be cautious in advocating for automation and in criticizing automation taxes. But, we must also avoid apologism. Maybe, given our propensities to ignore the claims of the least advantaged, we're pragmatically justified in using taxes to slow technological progress, despite the harmful economic consequences of doing so. But, if that's the case, then we should recognize this decision for what it is, namely a

compromise that we must endure because of our unwillingness to do what's really required of us.¹³

Author Details

Tom Parr

Department of Politics and International Studies

University of Warwick

t.parr@warwick.ac.uk

References

Abbott, Ryan and Bret Bogenschneider, 'Should Robots Pay Taxes? Tax Policy in the Age of Automation', *Harvard Law and Policy Review*, 12 (2018), 145-175.

¹³ Work on this article received funding from the European Union's Horizon 2020 research and innovation program under Marie Skłodowska-Curie grant agreement 890434. For helpful feedback on various of the ideas presented in this paper, I thank audiences at the Australian National University, the City University of New York, Harvard University, the London School of Economics, Princeton University, Queen's University, the University of Auckland, and the University of Bristol. For helpful discussions and/or written comments, I thank David Axelsen, Ayeek Bhattacharya, Christian Barry, Geoff Brennan, Matthew Clayton, Shuk Ying Chan, James Christensen, Tim Fowler, Bob Goodin, Alex Gourevitch, Johannes Kniess, Rahul Kumar, Matt Landauer, Kian Mintz-Woo, Gina Schouten, Kai Spiekermann, Lucas Stanczyk, Annie Stilz, Christine Sypnowich, Nelson Tebbe, Laura Valentini, Andrew Walton, Martin Wilkinson, Andrew Williams, and Peter Wilson. Special thanks to Will Abel.

Abel, Will, Elizabeth Kahn, Tom Parr, and Andrew Walton, *Introducing Political Philosophy: A Policy-Driven Approach* (Oxford: Oxford University Press, 2021).

Acemoglu, Daron, 'Technical Change, Inequality, and the Labor Market', *Journal of Economic Literature*, 40 (2002), 7-72.

Acemoglu, Daron and David Autor, 'Tasks and Technologies: Implications for Employment and Earning' in Orley Ashenfelter and David Card (eds), *Handbook of Labor Economics*: Volume 4 (Amsterdam: Elsevier, 2011), 1043-171.

Acemoglu, Daron, Andrea Manera, and Pascual Restrepo, 'Tax, Automation, and the Future of Labor', September 29 2020, available at <https://workofthefuture.mit.edu/wp-content/uploads/2020/10/2020-Research-Brief-Acemoglu-Manera-Restrepo.pdf>.

Acemoglu, Daron, Andrea Manera, and Pascual Restrepo, 'Does the US Tax Code Favor Automation?', NBER Working Paper 27052.

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

Arulampalam, Wiji, 'Is Unemployment Really Scarring? Effects of Unemployment Experiences on Wages', *The Economic Journal*, 111 (2001), 585-606.

Atkinson, Anthony, *Inequality – What Can Be Done?* (Cambridge: Harvard University Press, 2015).

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

Autor, David, Mindell, David and Elisabeth Reynolds, 'The Work of the Future: Building Better Jobs in an Age of Intelligent Machines', November 17 2020, available at <https://workofthefuture.mit.edu/wp-content/uploads/2020/11/2020-Final-Report2.pdf>.

Blank, Rebecca and David Card, *Finding Jobs: Work and Welfare Reform* (New York: Russell Sage Foundation, 2000).

Bogenschneder, Bret, 'Will Robots Agree to Pay Taxes? Further Tax Implications of Advanced AI', *North Carolina Journal of Law & Technology*, 22 (2020), 1-56.

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

Ci, Jiwei, 'Agency and Other Stakes of Poverty', *Journal of Political Philosophy*, 21 (2014), 125-50.

Clayton, Matthew and Zofia Stemplowska, 'Dignified Morality', *Jurisprudence*, 6 (2016), 309-26.

Crosby, Faye, *Spouse, Parent, Worker: On Gender and Multiple Roles* (New Haven: Yale University Press, 1987).

Dworkin, Ronald *Justice for Hedgehogs* (Cambridge: Harvard University Press, 2011).

Estlund, Cynthia, 'What Should We Do After Work? Automation and Employment Law', *The Yale Law Review*, 128 (2018), 254-326.

Frey, Carl Benedikt and Michael A. Osborne, 'The Future of Employment: How Susceptible are Jobs to Computerization?', *Technological Forecasting and Social Change*, 114 (2017), 254-80.

Goodin, Robert E., 'Theories of Compensation', *Oxford Journal of Legal Studies*, 9 (1989), 56-75.

Gottschalk, Peter and Richard Freeman (eds), *Generating Jobs: How to Increase Demand for Less-Skilled Workers* (New York: Russell Sage Foundation, 1998).

Hemel, Daniel, 'Does the Tax Code Favor Robots?', *Ohio State Technology Law Journal*, 16 (2020), 219-243.

Hicks, John, 'The Foundations of Welfare Economics', *Economic Journal*, 49 (1939), 696-712.

Jütten, Timo, 'Dignity, Esteem, and Social Contribution: A Recognition-Theoretical View', *Journal of Political Philosophy*, 25 (2017), 259-80.

Kaldor, Nicolas, 'Welfare Propositions in Economics and Interpersonal Comparisons of Utility', *Economic Journal*, 49 (1939), 549-52.

Keynes, John Maynard *A Tract on Monetary Reform* (London: MacMillan and Co, 1923).

Keynes, John Maynard, 'The General Theory of Employment', *Quarterly Journal of Economics*, 51 (1937), 209-23.

Lamont, Michelle, *The Dignity of Working Men: Morality and the Boundaries of Race, Class, and Immigration* (New York: Russell Sage, 2000).

- Mazur, Orly, 'Taxing the Robots', *Pepperdine Law Review*, 46 (2019), 277-330.
- McCredie, Bronwyn, Kerrie Sadiq, and Larelle Chapple, 'Navigating the Fourth Industrial Revolution: Taxing Automation for Fiscal Sustainability', *Australian Journal of Management*, 44 (2019), 648-64.
- McGoogan, Cara, 'South Korea Introduces World's First 'Robot Tax'', *The Telegraph*, August 9 2017, at <http://www.telegraph.co.uk/technology/2017/08/09/south-korea-introduces-worlds-first-robot-tax/>.
- Nagel, Thomas, *Mortal Questions* (Cambridge: Cambridge University Press, 1979).
- Nozick, Robert, *Anarchy, State, and Utopia* (New York: Basic Books, 1974).
- Parr, Tom, 'The Significance of Employment' (unpublished manuscript).
- Pigou, A. C., *The Economics of Welfare* (London: MacMillan, 1960 [1932]).
- Piketty, Thomas, *Capital in the Twenty-First Century* (Cambridge: Harvard University Press, 2014), 515-39.
- Rawls, John, *A Theory of Justice: Revised Edition* (Cambridge: Harvard University Press, 1999).
- Reuters, 'European Parliament Calls for Robot Law, Rejects Robot Tax', February 16th 2017, at <https://www.reuters.com/article/us-europe-robots-lawmaking/european-parliament-calls-for-robot-law-rejects-robot-tax-idUSKBN15V2KM>.
- Scanlon, T. M., *Why Does Inequality Matter?* (Oxford: Oxford University Press, 2018).

Shelby, Tommie, *Dark Ghettos: Injustice, Dissent, and Reform* (Cambridge: Harvard University Press, 2016).

Soled, Jay A. and Kathleen DeLaney Thomas, 'Automation and the Income Tax', *Columbia Journal of Tax Law*, 10 (2018), 1-48.

Wilke, Jeff, 'Small Business Success in Challenging Times', July 21 2020, available at <https://www.aboutamazon.com/news/small-business/small-business-success-in-challenging-times>.

World Bank, *World Development Report 2019: The Changing Nature of Work* (Washington: World Bank, 2019).