

Will life be worth living in a world without work? Technological Unemployment and the Meaning of Life

By John Danaher,
forthcoming in *Science and Engineering Ethics*

Abstract

Suppose we are about to enter an era of increasing technological unemployment. What implications does this have for society? Two distinct ethical/social issues would seem to arise. The first is one of distributive justice: how will the (presumed) efficiency gains from automated labour be distributed through society? The second is one of personal fulfillment and meaning: if people no longer have to work, what will they do with their lives? In this article, I set aside the first issue and focus on the second. In doing so, I make three arguments. First, I argue that there are good reasons to embrace non-work and that these reasons become more compelling in an era of technological unemployment. Second, I argue that the technological advances that make widespread technological unemployment possible could still threaten or undermine human flourishing and meaning, especially if (as is to be expected) they do not remain confined to the economic sphere. And third, I argue that this threat could be contained if we adopt an *integrative* approach to our relationship with technology. In advancing these arguments, I draw on three distinct literatures: (i) the literature on technological unemployment and workplace automation; (ii) the antiwork critique — which I argue gives reasons to embrace technological unemployment; and (iii) the philosophical debate about the conditions for meaning in life — which I argue gives reasons for concern.

Keywords: Automation; Technological Unemployment; Antiwork; Freedom; Egalitarianism; Meaning of Life

1. Introduction

Suppose the most extreme predictions regarding technological unemployment come to fruition: automating technologies takeover almost all forms of human labour. But suppose also that everyone shares in the productive gains. What happens then? Will our lives be filled with meaningful activity, or will we be merely passive recipients of the benefits of technology? In this article, I argue that there is cause for optimism, but only if we prioritise and develop the right kinds of technology and we relate to these technologies in the right way.

The main argument can be simply stated: Although there are reasons to worry about the societal impacts of technological unemployment, there are also reasons to embrace it. The so-called *antiwork critique* allows us to see the myriad ways in which our lives could be improved by substituting robotic labour for human labour: we would be free to pursue our own conception of the good life; our health and well-being could be enhanced; and we could rid ourselves of the drudgery, coercion and degradation of many forms of work. At the same time, the philosophical understanding of what it takes to live a meaningful life shows us the various ways in which the automation of labour could rob us of meaningfulness. This is true across a number of plausible theories of meaning, and should give us some reason for concern. But it may be possible to mitigate these concerns by harnessing technology in the right way. In particular, if we pursued increased *integration* with technology, not increased *externalisation*.

The argument builds over four main sections. First, I clarify the concept of work and explain why widespread technological unemployment is something we need to take seriously. Second, I introduce the so-called antiwork critique, noting how it highlights both the intrinsic badness and opportunity costs of work, and arguing that it is considerably more persuasive in an era of rapid technological advance. Third, I investigate the philosophical literature on meaning in life and make the case for thinking that technological unemployment threatens meaning. And fourth, I argue that this leads to a (mild) endorsement of an integrationist approach to technology.

2. Will Technological Unemployment Happen?

This article is about work and the threat of technological unemployment. It would help if we had a firmer grasp on these two concepts. I start with “work”. It is notoriously difficult to define “work” in a manner that both covers all the phenomena of interest, and avoids begging the question as to its desirability (Frayne 2015, 17-23). Consider Bertrand Russell’s famous and oft-quoted definition (2004):

“Work is of two kinds: first, altering the position of matter at or near the earth’s surface relatively to other such matter; second, telling other people to do so. The first kind is unpleasant and ill paid; the second is pleasant and highly paid.”

Although there is a delightful insouciance to this, it seems both over-inclusive and value-laden. Russell is appealing to a scientific definition in describing work of the first kind, and he is assuming that such work is always unpleasant. He is appealing to a purely managerial/consultative conception of other types of work and assuming them all to be well paid. This seems both unhelpful and untrue.

How can we avoid these definitional problems? There is no ideal solution, but I propose the following as a plausible definition of work, which should suffice for the purposes of the present discussion:

Work: The performance of some act or skill (cognitive, emotional, physical etc.) in return for economic reward, or in the ultimate hope of receiving some such reward.

This definition is quite broad. It covers a range of potential activities: from the hard labour of the farm worker, to the pencil-pushing of the accountant and everything in between. It also covers a wide range of potential rewards: from traditional wages and salaries to other reciprocal benefits. It explicitly includes forms of “unpaid employment”. Thus, for example, entrepreneurial work, unpaid internships and apprenticeships are included within my definition because, although they are not done in return for immediate economic reward, they are done in the hope of ultimately receiving some such reward.

Despite this broadness, I think the definition avoids being overly-inclusive because it links the performance of the skill to the receipt of some sort of economic reward. Thus, it avoids classifying everything we do as work. There are two potentially controversial consequences of this. The first is that the definition excludes certain types of care-work or domestic work, traditionally performed by women, that have gone without economic reward. It is not my intention to downplay the value of such work, or to dispute its classification as ‘work’, but I think the exclusion is justified because the present inquiry is concerned with how technological unemployment improves or disimproves things with respect to the current status quo. The lack of economic reward for domestic labour is part of the current status quo and so not directly relevant to the argument I wish to make (though the impact of technology outside the labour market will not be ignored). It may also exclude historical practices like subsistence farming in feudal societies or slavery that many would classify as ‘work’ (perhaps the exemplary forms of objectionable work).¹ These practices could be included if you loosen up the definition of the instrumental gains that are tied to the performance of the skill. But, again, seeking to include them may not be necessary given the focus of the present inquiry. We are concerned with the forms of work that dominate in contemporary capitalistic societies and the potential shift away from these forms of work because of technological advance. In this respect, I think the proposed definition does capture the core phenomenon of interest and it doesn’t beg the question by simply assuming that such work is, by definition, “bad”. The definition is silent on this issue.

So much for work, what about technological unemployment? I define that as the replacement of human workers, engaging in work as defined above, by technological alternatives (machines, computer programs, robots and so forth). Such technological unemployment is a long-standing feature of the economy. Famously, machines have been used to supplement or replace agricultural or manufacturing labour. More recently, with the growth of artificial intelligence and robotics, they have begun to replace cognitive and emotional labour. This leads some economists to refer to this as “the second machine age” (Brynjolfsson and McAfee 2014).

In discussing technological unemployment, so-defined, one must distinguish between short-term and long-term versions of the phenomenon. To date, technological unemployment has not caused any long-term economy-wide crisis. Displaced workers

¹ I am indebted to an anonymous reviewer for drawing this problem to my attention.

(or, at least, future generations of workers) have always found other forms of work (Autor 2015a). Thus, most forms of historical technological unemployment have been short-term. The phenomenon of interest in this article is the long-term structural version of technological unemployment, i.e. a version in which displaced workers, and subsequent generations of workers, cannot find alternative forms of work.

This type of technological unemployment has become a serious matter of concern in recent years.² The goal of this article is not to defend the credibility of this concern. Rather, the goal is to assume that technological unemployment is going to happen and to consider its social and ethical implications. Nevertheless, it is worth getting some sense of the case in favour of technological unemployment. Doing so serves two purposes: It helps to justify the present inquiry, and it highlights the various technological trends that might facilitate such structural unemployment. This is particularly important since I will be arguing in section 4 that these trends will impact on our non-work lives as well.

The case for technological unemployment looks something like this:

(1) If technology is replacing more and more forms of human labour, and if there are fewer and fewer alternative forms of work for humans to go to, then there will be technological unemployment.

(2) Technology is replacing more and more forms of human labour, and is doing so in a way that results in fewer and fewer forms of alternative work for humans.

(3) Therefore, there will be technological unemployment.

The first premise is pretty uncontroversial. It states an obvious causal connection between replacement by technology, a narrowing field of employment opportunities, and the increase in unemployment. The second premise is problematic. It has two

² A number of books, some quite alarmist and pessimistic (Keen 2015; Carr 2015), some more cautious and optimistic (Brynjolfsson and McAfee 2012 & 2014; Ford 2009 & 2015; Pistono 2012; Cowen 2013; Kaplan 2015; Rifkin 1997 & 2014), have been published arguing that ours is an age of increasing technological unemployment. These books have been complemented by research papers highlighting the rise of automation and the increasing share of income being taken by capital in Western economies (Frey and Osborne 2013; Fleck, Glaser and Sprague 2011; ILO 2013; Pratt 2015; Sachs, Benzell & LaGarda 2015). These have in turn been complemented by the work of a number of leading journalists and economic opinion writers (Packer 2013; Krugman 2012 & 2013).

components. Its first component can be defended by appealing to enumerative examples of technologies that are replacing human labour (Brynjolfsson & McAfee 2014; Ford 2015; and Kaplan 2015). It can also be supported by appealing to some trends in the key economic indicators that suggest that technology is responsible for an increasing amount of economic output (ILO 2013; Ford 2015).

The main problem with premise (2) is its second component which commits the Luddite fallacy³ (Brynjolfsson & McAfee 2014). This is the fallacy of believing that replacement by technology necessarily narrows the field of employment opportunities. This is a fallacy for both empirical and theoretical reasons (Autor 2015a, 2015b & Forthcoming). Since the dawn of the industrial revolution machines have been replacing workers and yet workers have always found somewhere else to go. Defenders of this view argue that machines complement human labour, they never completely replace it (Autor 2015a). The modern case for technological unemployment must overcome this Luddite fallacy. And there are a number of plausible rebuttals. Four, in particular, stand out among defenders of technological unemployment (Brynjolfsson and McAfee 2014).⁴

First, they point to the *inelastic demand problem*. There may be limits to how much of particular good or service is ever demanded and this could have a negative impact on jobs no matter how much the machines complement human workers. A clear example of this is the demand for labour in agriculture and manufacturing. Increased machine productivity has not led to increased demand for human labour in those industries. Even those who are sceptical of technological unemployment concede this example, but argue that it is merely one among many and that the ultimate impact of technology depends greatly on the particular labour market (Autor 2015a).

Second, they point to the *outpacing problem*. The Luddite fallacy assumes that people will always be able to retrain and reskill themselves (or that future generations will be able to do so) at a rate that cannot be matched by improvements in technology. But this may no longer be true. Many now argue that improvements in technology are

³ This can also be referred to as the lump of labour fallacy, i.e. the fallacy of believing that there is a fixed lump of labour out there to be distributed among human workers.

⁴ Brynjolfsson and McAfee also mention a fifth: the possible creation of android robots, i.e. perfect technological replicas of humans, only cheaper, more compliant and more efficient. This might be an apotheosis of the current technological trends. But some argue that it is a mistake to focus on human-like machines; it is the fact that machines are not human-like that makes them a displacement threat to human workers (Kaplan 2015).

exponential (Kurzweil 2006; Ridley 2011; Agar 2015). If it takes several years to retrain a worker (or 25-30 years to educate them from scratch), that may no longer be enough to keep pace.

This is connected to a third problem, which we can call the *historical data problem*. This suggests that any empirical support for the Luddite fallacy could stem entirely from observing the relatively linear portion of an exponential growth curve in technological advancement. In other words, the reason why we haven't yet seen long-term structural unemployment is because we have yet to see what happens when technological advances really start to take off (Brynjolfsson & McAfee 2014).

Fourth, and perhaps most importantly for this article, they point to the *winner takes all problem*. There are two key trends in the modern information economy. The first is that an increasing number of goods and services are being *digitised* (with the advent of 3D printing, this now includes physical goods). This allows for replication at near zero marginal cost (Rifkin 2014). A second key trend is the existence of *globalised* networks for the distribution of goods and services. This is obviously true of digital goods and services, which can be distributed via the internet. But it is also increasingly true of non-digital goods, which can rely on vastly improved communication and transport networks for near-global distribution. These two trends have led to more and more “winner takes all” or “superstar” markets in which being the second (or third or fourth...) best provider of a good or service is not enough: the income tends to flow to one dominant participant (Brynjolfsson and McAfee 2014; Keen 2015). Services like Facebook, Youtube, Google and Amazon dominate particular markets thanks to globalised networks and cheap marginal costs. The fact that the existing infrastructure makes winner takes all markets more common has devastating implications for long-term employment. If it takes less labour input to capture an entire market — even a new one — then new markets won't translate into increased levels of employment. The fact that technology allows for such superstar markets has important implications for what we do outside of work. I return to this point in section 4.

None of these responses is watertight. After the deluge of pro-technological unemployment pieces in recent years, a backlash has emerged. Several authors argue that the current and future impact of technology on work is being overstated (Autor 2015a, 2015b and forthcoming; Mokyr, Vickers and Ziebarth 2015) and that there are

other, more pressing, threats to employment (Denning 2015). More needs to be said, but saying it lies beyond the scope of this article. This brief summary should, nevertheless, enable the reader to see why people are genuinely worried about the prospect of technological unemployment, which technological trends are fueling this worry, and hence why it is worth seriously considering its social and ethical implications. That is what the remainder of this article will do.

3. The Case for Optimism in the face of Automation

If we assume that technological unemployment is a likely possibility, two major social and ethical questions arise: (i) how will people access the essential and non-essential goods and services that currently require an income from work? and (ii) what will people do with their time now that they are not working?

The first question points to a *distributional problem*. Technological unemployment implies that a greater share of income will go to capital (and the owners of capital) than to labour. One presumes that the owners of capital will be few while the displaced workers will be many -- this is confirmed by recent analyses (Piketty 2014; Atkinson 2015; Autor 2015a). This implies greater income and social inequality. This can only be resolved if there is some significant redistribution of wealth from the capital owners to the displaced workers (or some other more drastic reorganisation of the economy). Many technologists and futurists have suggested that an unconditional basic income guarantee could resolve these problems.⁵ I take no view on the matter here. Instead, I simply assume that the distributional problem can be solved: the technologically unemployed future will be one of *abundance*, not deprivation. This is a significant assumption. It means that everything I say below is conditional on the assumption that people won't be suffering from great hardship due to the lack of an income.

The second question points to a *problem in personal axiology*, i.e. in what makes life fulfilling and worthwhile for the person living it. There is a widespread view that

⁵ This is defended in Ford (2015) and Brynjolfsson & McAfee (2014), though the latter refer to the reverse income tax, which is effectively the same policy. Technologist and futurist James Hughes believes that technological unemployment provides a strategic opening for proponents of the basic income (2014). For further discussion of the basic income proposal, see Widerquist (2013); Widerquist et al (2013); van Parijs (1995); Ackerman et al (2006).

work is virtuous (Weeks 2011; Srnicek and Williams 2015) and the absence of paid employment could lead to vicious forms idleness, boredom, and depression. Working for a living confers a certain level of well-being and individual flourishing that would be absent if one did not work for a living.⁶ For anyone who accepts this view, there is much to worry about when it comes to the prospect of technological unemployment. There is a danger that in robbing us of the dignity of work, increasing levels of automation could reduce our overall levels of flourishing. The goal for the remainder of this article is to develop and evaluate the arguments associated with this second problem.

I start by defending an optimistic view which draws upon the existing antiwork critique of capitalist work culture.⁷ There are two main families of arguments that form the backbone of this antiwork critique. The first are what I shall call ‘work is bad’ arguments. These arguments claim that technological unemployment is to be welcomed because it will take away something that is bad for us, both individually and socially. The second are what I shall call ‘opportunity cost’ arguments. These arguments claim that technological unemployment is to be welcomed, even if work is not bad, because non-work is simply better. I argue that both sets of arguments become more persuasive in an era of technological unemployment.

The badness of work is repeatedly highlighted by antiwork theorists. Bob Black,⁸ in his classic call to arms ‘The Abolition of Work’, argues that:

“Work is the source of nearly all the misery in the world. Almost any evil you’d care to name comes from working or from living in a world designed for work. In order to stop suffering, we have to stop working.”

(Black 1986)

⁶ This view is implicitly or explicitly evoked in the automation-related work of Carr (2015), Brynjolfsson and McAfee (2014), and Ford (2015).

⁷ There is an extensive literature on this topic, examples include Lafargue (1883); Black (1986); Gorz (2011); Woodcock (1944); Graeber (2013); Crary (2014); Russell (2004); Levine (1995); Maskivker (2010); Widerquist (2013); Weeks (2011); Schwartz (1982); Srnicek and Williams (2015).

⁸ Black’s work is extreme and polemical but is influential in the antiwork movement. Frayne (2015, 206-207) notes that members of the Idler’s Alliance (an organisation seeking to resist work) ‘repeatedly recommended’ Black’s essay to him.

Similarly, Bertrand Russell suggests that ‘a great deal of harm’ is done by the belief that work is virtuous (2004); and Kathi Weeks (2011), in her Marxist-feminist analysis of work, argues that there is something puzzling about our commitment to work when it seems true that work itself is so bad both in terms of its quality and quantity.

But what exactly does the badness of work consist in? There are two ways to answer to that question. The first is to highlight certain specific and contingent properties of different classes of work. Thus, it can be argued that few people are happy with their work;⁹ that many forms of work are degrading and humiliating (e.g. certain types of erotic and care work); that some are correlated with high rates of mental illness, depression and suicide (e.g. social work and legal work); that some expose workers to unacceptable physical risks (e.g. sweatshop labour); that workplaces are sometimes rife with bullying, intimidation and sexual harassment; that work often leaves one exhausted and deflated; and so on. The exhaustive enumeration of such bad-making properties would, no doubt, be instructive, but also quite limited. Opponents of the antiwork view could just as easily highlight good-making properties of work, such as the sense of accomplishment and social status that comes with it (Gheaus and Herzog, forthcoming), and the fact that working enables one to access other social goods and services.

The problem is that the features that are typically highlighted are contingent upon particular forms or classes of work. Thus, even if everyone agrees that cleaning toilets is not the most glamorous way to make a living, they can point out that cleaning toilets is just one form of work among many. Indeed, one reason to embrace technological displacement in the workplace could be that it frees up workers to pursue more creative and meaningful forms of work. Thus, if we focus solely on contingent bad-making features of certain forms of work, we may end up welcoming some technological displacement, but maintaining our overall commitment to the value of work and the work ethic. I want to defend the more radical view that we should abandon that commitment.

This is where the second type of argument comes in. This one highlights bad-making features of work that are intrinsic to the contemporary economic-political structure of work. That structure includes the fact that work is something that one *must*

⁹ In support of this one could cite the 2013 Gallup survey suggesting that only 13% of workers worldwide were engaged by their work. See <http://www.gallup.com/poll/165269/worldwide-employees-engaged-work.aspx> (accessed 26/10/15).

do in order to access the basic necessities and luxuries that make life possible and worth living. In other words, in the contemporary economy, work is *intrinsically compulsory* in nature. This is a controversial claim and needs to be finessed. It is not that we are wage-slaves. We are not physically or legally obliged to work. Rather, work (and the pursuit of work) is a practical and economic necessity, even in the absence of physical or legal obligation.

Some might challenge this on the grounds that contemporary welfare states obviate the need for work. But such an argument would ignore the fact that most forms of welfare, in most states, are both insufficient for survival, and explicitly tied to the ability to work. One must usually prove either an active willingness to seek work, or a physical/mental inability to work, in order to access welfare entitlements.¹⁰ Thus, even in a relatively generous welfare state, one cannot escape the work ethic except through illness, old-age and death. It is in this way that we are all compelled to pursue work and to participate in a culture that glorifies work. This has at least two bad-making properties that could be eliminated in a postwork future.

The first is that the compulsory nature of work leads to an impoverished system of egalitarian justice (Levine 1995). One of the central precepts of liberal egalitarianism is that the state should be neutral with respect to its citizens' conception of the good life.¹¹ It should not promote or force any particular conception of the good life on its citizens. It should work to tolerate and facilitate people in their pursuit of different conceptions of the good. Of course, it can only do this to the extent that a person's conception of the good does not unjustly or unfairly deprive anyone else of their conception of the good. The key move for the antiwork proponent is to argue that non-work is a model of human flourishing that the neutral state should tolerate. And the problem with the compulsory nature of work is that it fails to tolerate and facilitate such alternative conceptions of the good life.

¹⁰ This ignores the recent social push toward, and limited experimentation with, unconditional basic income schemes. I ignore this for two reasons: (i) the rollout of such schemes is minimal and typically not sufficient to meet basic needs; and (ii) the increase in such schemes is arguably being driven by concerns about structural unemployment.

¹¹ Appealing to the goodness of the neutrality of the state with respect to citizens' choice of the good life does not imply any particular view of whether non-work is better than work. It merely claims that the state should not impose a choice on its citizens. This leaves open the door to arguing that non-work would be the better choice. I am indebted to an anonymous reviewer for encouraging me to make this clarification.

This is sometimes defended on the grounds that the state *cannot* facilitate such lifestyle choices: doing so would result in ‘lazy’ people being supported by other hard-working, economically productive individuals (van Parijs 1995). That would lead to injustice. But this is where anti-work arguments start to make more sense in an era of technological unemployment. If we are about to enter an era of technological unemployment, this allows for a re-ordering of the economic-political reality of work: machines can takeover the economically productive forms of labour and, provided the gains from this are appropriately redistributed, the ‘lazy’ lifestyle can be tolerated without a corresponding injustice to others.

A second, and perhaps more persuasive, argument can be found in the work of Julia Maskivker (2010). She argues that compulsory work is bad because it undermines a core value of the liberal democratic state: freedom/autonomy.¹² Many people would agree that freedom is an essential part of the flourishing life, and compulsory work certainly seems to undermine it. The strength of Maskivker’s analysis lies in identifying three specific ways in which this happens. They are:

Freedom-undermining Properties of Work: Compulsory work (i) limits our ability to choose how to make use of our time; (ii) limits our ability to be the authors of our own lives; and/or (iii) involves exploitative/coercive offers.

Each of these bad-making properties deserves some scrutiny.

The problem with limiting how we can choose to spend our time is best understood in relation to the right to control self-ownership (Widerquist 2013). In the Lockean view, individuals have a right of ownership over their bodies and the fruits of their labour. The problem is that there are resources you need for effective self-ownership. For example, people may need an adequate education and access to healthcare. Time is also an essential resource for self-ownership (Maskivker 2010). In fact, time is probably the ultimate resource (Zeckhauser 1973). Other skills and abilities only really have value when we have the time to exercise them. Furthermore, time is a peculiarly non-manipulable resource. There is a limited amount of time in which we get to act out our lives. This makes it all the more important for people to have access to

¹² These can be understood as equivalent or distinct concepts. Gerald Dworkin (1981) argues, for instance, that freedom is local concept that applies to particular decisions, whereas autonomy is a global concept applying across a larger swathe of decision points.

time. The problem, as Weeks (2011) and others point out, is that work tends to monopolise our time. The modern knowledge-oriented workplace is particularly insidious in this regard.¹³ Modern communications technologies have broken down the barriers between work and leisure (Crary 2014). We are encouraged to use our time more productively, but also to be available to our workplaces at more times of the day. The death of the regular 9-to-5 workday has, if anything, encouraged work to monopolise more of our time. We have flexible working hours and our work is more outcome-driven, but the marketplaces are open 24/7 and they demand more outcomes from us. The result is an infiltration of work into every hour of the day. Dissolving this monopolisation of time is something we could look forward to in a more fully-automated economy.

The way in which work limits our authorship of our own lives, and the badness this entails, is also linked to the ideal of self-ownership. To understand the idea, we need to think about the individual who truly enjoys their work. Such an individual derives great personal satisfaction and meaning from their work. In this, they seem to be truly owning the narrative of their lives. You could consequently argue that work has no bad-making properties for such an individual. But there are reasons to doubt this. Maskivker uses the analogy between starving and fasting to make the point. When a person is starving or fasting, the physical results are often the same: their bodies are being deprived of essential nutrients. But there is something morally distinct about the two cases. The person who chooses to fast has authorship over their lives; the person who is starving is having their story written by someone else. When it comes to work, there is a sense in which we are all starving not fasting. We may enjoy it, embrace it and endorse it, but at the end of the day we have to do it. We do not have the true authorial control of the faster. This is another reason to embrace technological unemployment. Freeing us from the necessity of work will remove a barrier to authorial control.

This brings us to the final bad-making property of work. It is related to the previous properties but can be treated separately. It is that work, in the modern economy, involves a coercive/exploitative offer. You are being asked to work ‘or else’, i.e. work or forgo a number of opportunities or necessities. Thus, even if work ends up being beneficial, it is not treating the worker as a truly autonomous agent. Some might

¹³ The way in which industrial capitalism polices and manages our times has been long discussed in the antiwork literature, see Woodcock (1944)

dispute the coercive or exploitative of the work-offer by arguing that coercive offers always require a threat to make *you worse off than you might otherwise have been* and it is not clear that the work-offer involves such a threat (Wertheimer 1988). The idea here would be that work is essential for economic productivity and that economic productivity raises societal well-being. In a world without compulsory work we would be much worse off. Thus, there is no sense in which the work-offer is threatening to make you worse off than you might otherwise have been.

But, again, this is one place where technological unemployment makes a real difference to the strength of the antiwork position. It might be true in certain social orders that you would be worse off without work, but in a society in which machines can takeover most forms of productive labour, this is much less clear. In such a society, work seems like, at best, an unnecessary extravagance and, at worst, a form of slow torture. If the gains of machine labour can be appropriately redistributed, then we are not worse off for not working. And, indeed, if work retains its compulsory nature in such a society, it would become truly coercive, within the terms set out by the likes Wertheimer (1988). In other words, whatever you think about this bad-making property of work in contemporary societies, it would be particularly acute in a world of rampant technological unemployment.

These ‘work is bad’ arguments can be complemented by a set of ‘opportunity cost’ arguments. These arguments accept that work might have many good-making properties but insist that non-work is better. The opportunity cost arguments have been implicit in the discussion thus far. They can now be made explicit. For starters, it is worth pointing out that one’s choice of work is always limited by what kinds of activities are economically viable. And the kinds of activities that are economically viable are dependent on where wealth is located and how that wealth is disposed (Graeber 2011). There is no guarantee that distribution via the labour market allow workers to spend their lives in the pursuits they find most fulfilling. If we could break the link between income and how we choose to spend our time, then we could spend our time in the manner that best suits our skills, desires and aptitudes. Thus, the biased nature of the market gives us some reason to think that non-work is preferable to work.

There are other instrumental and intrinsic reasons for thinking that non-work is better. Russell (2004), for instance, argues that the leisure classes (i.e. those free from

the necessity of work) have always been responsible for the scientific, cultural, and political developments that form the backbone of our civilisation. Consequently, he believes that freeing up more people into the leisure class will make for a better world. Bob Black takes a slightly different tack. He argues that the 'ludic life' (i.e. the playful life, free from work) is intrinsically better than the working one, and believes that it could lead to "a collective adventure in generalized joy and freely interdependent exuberance" (Black 1986).

These 'opportunity cost' arguments can be resisted on the grounds that even if the ludic, leisurely life is better than the working one, work itself is necessary in order to secure greater goods (e.g. supply of food and healthcare for all, money to pay for leisure activities). Or, alternatively, on the grounds that they ignore the opportunity costs of non-work itself. Income is not the only good we derive from work in the modern world. As Gheaus and Herzog (forthcoming) have recently argued, there are at least four non-monetary goods associated with paid work: (i) excellence in the cultivation of our skills; (ii) contribution to society; (iii) a sense of community and (iv) social status. People want to be good at what they do; they want to make a difference to the world in which they live; they want to engage in collective enterprises; and they want to be recognised and validated in the eyes of their social peers. Paid work is the main forum in which they attain these four goods. Are we going to be cutting people off from these sources of meaning and flourishing in a post-work world?

I would argue that both forms of resistance are much less compelling if we are living through an era of technological unemployment. The prospect of near ubiquitous automating technologies seems to break the allegedly necessary link between work and these goods. Gheaus and Herzog themselves acknowledge that paid work is a privileged context for the attainment of the four goods they identify partly because we have to spend so much of our time doing it. It thus becomes the only real outlet we have for achieving excellence, social contribution, community and social status. But there is nothing intrinsic to work as I define it in this paper (performance of a skill for economic reward) that makes it the only forum for achieving those goods. People can, and do, achieve excellence, social contribution, community and status in leisure activities, voluntary charitable activities and hobbies. They tend not to because they are forced to spend their time in economically productive forms of work. If technology can take over

the economically productive activities, then the instrumental and intrinsic virtues of these non-work contexts can flourish.

That said, there is a problem with the antiwork position defended thus far. The combination of the ‘work is bad’ and ‘opportunity cost’ arguments do provide us with a robust *prima facie* reason to welcome technological unemployment. But there is something missing. In their present form, these arguments are undertheorised. In order to make them fully persuasive we need to consider more closely what makes for the good life and how exactly the absence of work will facilitate that life. When we do this, we end up with a more nuanced and less optimistic view.

4. How Technological Advance could undermine the Good Life

In this section, I refine the anti-work position by considering in more detail what it takes to live the good life, and by asking the question: are we sure that removing the incentives and motivations provided by work will allow us to live that life? To answer that question, it is worth considering in more detail the typical philosophical accounts of what makes for a meaningful life. The theories settle into four main categories:¹⁴

Simple Subjectivist Theories: Hold that a life is meaningful to the extent that the individual living it experiences certain subjective states, typically conscious well-being and desire satisfaction.¹⁵

Simple Objective Theories: Hold that life is meaningful to the extent that the individual living it brings about certain objectively good or valuable states of affairs (Smuts 2013). These typically involve making the world a better place (morally speaking), producing some great scientific or intellectual discovery, or making great art (Metz 2010; 2013).¹⁶

Aim-Achievement Theories: Hold that a combination of subjective and objective states are needed in order to make life meaningful. The individual should set

¹⁴ This categorisation is consistent with those offered by Smuts (2013) and Metz (2013), but is taken directly from Nyholm and Campbell (2015).

¹⁵ Proponents of simple subjective theories include Richard Taylor (2008) though he later changed his view; AJ Ayer and Klemke (Nyholm & Campbell 2015 for more)

¹⁶ Though note that Metz’s own preferred view is more akin to fitting fulfillment.

themselves objective goals, bring them about through their actions, and have the feeling of satisfaction associated with doing this, i.e. that they have an ‘aim’ and that they ‘achieve’ this aim (Luper 2014).

Fitting-Fulfillment Theories: Hold that a combination of subjective and objective states are needed in order to make life meaningful. They are similar to the aim-achievement theories in that they insist on the individual being subjectively satisfied by bringing about some objective end. Where they differ is in thinking that the objective end must be one that is objectively good or objectively valuable, *i.e. fitting* (Wolf 2010; Metz 2010; Wielenberg 2005).

We can use these theories to flesh out the opportunity cost argument in favour of technological displacement. There are two points to be made. First, under a simple subjectivist theory, there is reason to think that technological unemployment could enhance the overall level of meaning in our lives, but only if we make use of the right kinds of technological advances. Second, in relation to all the remaining theories of meaning (objective and hybrid) there is reason to think that technological unemployment could undermine the overall level of meaning in our lives, but this impact could be lessened with the right kind of technological developments. These two arguments need to be unpacked.

4.1 - The Subjective Satisfaction of Non-Work

Let’s start with the simple subjectivist argument. Simple subjectivist theories lurk behind many of the anti-work arguments outlined above. The idea is that compulsory work takes us away from the things that we are really passionate about and that would confer upon us the most subjective satisfaction. Even if we are lucky enough to work at something we genuinely enjoy we must bear in mind that others are not so lucky. Things would be much better, subjectively speaking, if we could spend our time as we saw fit. The attraction is obvious. As I sit here writing this article, I often think to myself of other activities I could be performing which would be much more subjectively fulfilling. I could be out swimming or bike riding; I could be reading a book or working on that novel I have always wanted to write. But I cannot do these things because the productivist ethos of modern academia demands that I produce more peer-reviewed publications to pad out my CV.

This argument assumes that if we control our own time we will spend it in a way that induces the right subjective states. This assumption is cast into doubt by findings in psychology. Dan Gilbert's work on mis-wanting, for example, suggests that we often don't really know what makes us truly happy and often stumble upon happiness (Gilbert 2005; Gilbert et al 1998; Gilbert et al 2004). Findings like this can be exploited by critics of automation and technological unemployment. A recent example of this is Carr (2015). Carr contends that without the pressures and incentives of work we may live a life of listless and unsatisfied boredom.

In making this argument, Carr focuses on the role of 'flow states' in subjective well-being. Flow states are states of total immersion in a task or activity. They were first described by Mihaly Csikszentmihalyi (1990; 1997; 2007) and are characterised by strong positive affective experience (sometimes described as being rapturous or joyful). If any psychological state is a candidate for the subjective basis for meaning, a flow state would seem to be it. The problem is that Csikszentmihalyi's studies suggest that people are much more likely to enter flow states at work than they are at leisure. Csikszentmihalyi and colleagues have used experience sampling methods (2007) to interrupt people during either work or play, and get them to answer questions about their overall levels of personal satisfaction and well-being. These studies reveal that people are generally more focused, happier and more satisfied at work than at play. Indeed, they often report feeling anxious and bored outside of work when they are presumably free to engage in their preferred activities. All this despite the fact that they still claim, when asked, to prefer play to work. This creates the so-called 'paradox of work'.

Carr thinks the reason for this 'paradox' is obvious. Csikszentmihalyi's theory holds that entering into a flow state is a function of how difficult the task is and how much pressure is associated with it. A simple task with no associated pressure is unlikely to generate a flow state; the same goes for a task that is too difficult and comes with too much pressure. Moderate-to-high levels of pressure and difficulty are needed. Work is often an excellent way to provide the right kind of pressure and difficulty. If left to our own devices, we may not push ourselves hard enough. There is, consequently, reason to doubt whether, in the absence of work, we will spend our free time engaged in more subjectively fulfilling activities.

There are two problems with this argument. The first has to do with the empirical basis for the claim that without work we would be less likely to engage in flow-state-inducing activities. There is a limited utility to studies, such as Csikszentmihalyi's, which compare work and leisure in a world that is dominated by work. When I come home from a busy day at the office, I'm usually drained and lethargic. I'm often not physically able to engage in the kinds of activity I would prefer. I'm conscious of the fact that I need to recover before going back to the office again. In an era of rampant technological unemployment, in which the shadow of work is removed, things could be very different.¹⁷

The second problem with Carr's argument is that it makes a dubious assumption about the link between work and the kinds of pressures that are needed to induce flow states. It is paternalistic to assume that people will lack sufficient self-motivation if they are unemployed. This paternalism is behind much of the traditional ideological glorification of the work ethic. But worse than that, this argument ignores non-work related pressures and motivations that could be leveraged to similar effect. I can set difficult personal goals and publicly commit to them in a manner that will prove costly (e.g. in terms of reputational loss or damage) if I fail to achieve them. This could well provide the pressure needed for entering into flow states. Indeed, people do this all the time, for example, by committing to run marathons by raising funds for charity, or by committing to change their habits by entering into pacts or agreements with their friends. What's more, the argument ignores the ways in which modern technology can greatly assist in providing these alternative sources of pressure. Social networking and gamification apps, for instance, can provide the pressures and rewards needed to push ourselves out of our comfort zone and engage in the kinds of activity that are more likely to induce flow states.

So although the simple subjectivist argument in favour of technological unemployment can be naive, there are ways to correct for this naivety by employing the right kinds of social/technological support for leisure activities. But even this revised subjectivist argument should not be pushed too far. There are two reasons for this. First, purely subjectivist theories of meaning are implausible. The classic counterexample is the story of Sisyphus pushing his boulder up a hill for eternity. Suppose he derives

¹⁷ I would like to thank KS for helping me to make this argument.

immense personal satisfaction from getting it to the top and the satisfaction remains the same no matter how often he repeats the task. Here, Sisyphus is living a subjectively satisfying life, but there is nothing deeply meaningful about it. We need something more for that.

Second, and linked to this problem, although technology could be leveraged in ways that make us more likely to achieve flow states through our activities, there are also ways in which we could use technology to trick ourselves into such subjectively pleasurable states without any associated activity. ‘Wireheading’ technologies may be feasible in an era of technological employment (Pearce 1995). These would utilise brain implants to induce profound feelings of joy and satisfaction. The kinds of lives lived with such technological assistance might be incredibly subjectively satisfying, but they wouldn’t seem to embody the qualities of the good life as traditionally conceived. In short, if we are to think seriously about meaning and personal fulfillment in an age of technological unemployment, we probably need to take into consideration the link between our actions and the objective world, and how technology might mediate the relationship between our actions and the objective world.

4.2 - Could technology sever our link with objective sources of meaning?

This is something that the three remaining theories of meaning encourage us to do. Each of them stresses the need for individual activity and subjectivity to join-up with consequences in the external world. The simple objectivist theory demands that our actions be *causally responsible* for objectively good outcomes (Smuts 2013), irrespective of whether we appreciate or understand that our actions are responsible for those outcomes. The other two theories — aim-achievement and fitting-fulfillment — also demand that our actions be causally responsible for objective outcomes, but add to this the requirement that we be *subjectively aware* of this, and *satisfied* and *fulfilled* by our achievements. The aim-achievement theory, in the form defended by Steven Luper (2014), allows for objectively wicked/evil outcomes to facilitate a meaningful life, provided that the individual aimed for and achieved those outcomes. This strikes me, and many others (Wolf 2010; Metz 2010; and Wielenberg 2005), as unsatisfactory, so from here I will assume that the link must be with objectively good outcomes. That said, most of what I am about to say applies equally to the achievement of bad outcomes because the central question is whether it is more or less likely that we will be able to *achieve* outcomes in a world of rampant technological unemployment?

A naive argument claims that it is more likely. We see this at play in some of the antiwork positions outlined above. Take Russell's claim (2004) that the leisure class has, historically, been responsible for all that is good in the world (art; scientific discovery; moral good works; political development). This claim rests on the belief that in the absence of work we will be more likely to engage in the activities that produce objectively good outcomes. This is often tied to the claim that the market can distort or bias human activity towards outcomes that are profitable but not necessarily objectively good. Consequently, there is reason to think that technological unemployment, by freeing us from these distorting influences, will allow us to engage in the kinds of activities that are essential to the good life.

This argument is naive for at least two reasons. First, it is too dismissive of the potential for the market to direct human activity towards objectively valuable outcomes. Economists have long lauded the moral potential of the market. Markets can ensure efficient distribution of social goods; and they can increase the size of the social 'pie'. They can also incentivise the kinds of activity that are thought to be conducive to the good life. These include the production of art and intellectual discovery, both of which are subject to significant market forces in the modern world. If automating technology renders human contribution to such market-based activities unnecessary, then we may be robbed of something that is conducive to the good life.

But this objection is not entirely compelling since the antiwork camp will simply respond by saying that nonwork is better at allowing us to do these things. This is the second respect in which the antiwork argument is naive. It assumes that the kinds of technological advance that make widespread technological unemployment possible will occur in a vacuum — that the impact of automating technologies will be felt solely in our economic lives. This is highly unlikely. Indeed, we can already see ways in which automating and assistive technologies infiltrate our moral, artistic, intellectual and political lives. If this trend continues, and we rely on those technologies in these other domains, we could sever the necessary causal and mental link between our actions and the outcomes that are said to be constitutive of meaning.

Some concrete illustrations can be used to underscore this argument. Suppose the hope is that in the absence of work we will be free to use our creativity in pursuit of

intellectual discovery and moral betterment. Suppose the hope is that we will be able understand more deeply the mysteries of human biology, in particular the causes and progression of cancer. This hope may be forlorn. As with most contemporary biosciences, cancer pathology is increasingly reliant upon machine-learning algorithms to analyse large datasets and identify useful patterns in those datasets. These algorithms are now demonstrating creativity and reaching useful conclusions that are surprising to their original human creators. A good example of this is the use of the C-Path algorithm to improve our understanding of the prognosis of breast cancer. The algorithm discovered that better predictions could be made by analysing the structural tissue around cancerous cells rather than by analysing the cancerous cells themselves. This was not expected by the original researchers and programmers (Beck et al 2011). They may have created the technology, but they did not ultimately control the discovery. What is true of the study of cancer pathology is becoming true of science more generally. Science is increasingly a ‘big data’ enterprise, reliant on algorithmic, and other forms of automated assistance, to process large datasets and make useful inferences from those datasets. Humans are becoming increasingly irrelevant to the process of discovery.

But that’s just scientific discovery, what about other areas of activity that lead to valuable outcomes? Suppose, once you are freed from the shackles of work, you want to address some moral problem. Suppose you want to ensure that certain resources are distributed more fairly or more efficiently. You might think that this is a space in which your activity can be causally responsible for good outcomes. But this too is becoming less and less clear. For example, matching algorithms can and are being used to efficiently and fairly distribute resources. Kidney exchanges are a good example (Roth 2015; Vulkan et al 2013; and Sonmez & Unver 2013). No one doubts that it is objectively good that people are allowed to live longer through organ transplants. But there is also a widely acknowledged deficit in the number of willing (or deceased) donors. Furthermore, even when there is a willing donor, they are often not a match for their intended donee. Matching algorithms can help to solve this problem (Roth 2015; Sonmez & Unver 2013). Using large databases of patient and donor profiles, they can be used to link together matching pairs, thereby kickstarting long chains or cycles of organ donation.

What is true in relation to organ donation is (and will be) true of other moral enterprises. Energy and food distribution (in addition to production) can be improved through algorithmic and automated systems (Rifkin 2014). And many legal-bureaucratic enterprises can be too (Pasquale 2015). Governments are increasingly turning to such systems to perform tasks like identifying terrorist suspects, stolen identities, tax cheats, and likely repeat offenders. These systems can be tireless; and some argue (Zarsky 2012) that they are free from the distorting biases of human beings (though this is hotly contested). These systems often still involve human actors, and given the nature of democratic political processes probably always will, but the roles of the humans will become more passive, less creative and interactive. The machines start doing much of the work themselves; the logic of their decision-making becomes more and more opaque to those who interact with them.¹⁸ Thus, once again, the rise of automation reduces the space in which humans can engage in meaningful and fulfilling moral activities.

Now, you might dispute this characterisation. You might argue that there is still room for human input in all of these automated systems. For one thing, such systems would seem to require human designers, programmers and overseers; for another, humans would still have to contribute to the smooth functioning of such systems, e.g. by becoming kidney donors or by providing crucial data. Thus, individuals can be linked to objective outcomes in the appropriate manner. But there are three problems with this argument.

First, not everyone is going to be equipped or trained to design or program such systems. At the very least, there needs to be a massive shift in education and training so as to allow people to do so — mere freedom from work will not solve the problem. Furthermore, even if people do have the necessary training and skills, modern information technology and globalisation is such that relatively few people are needed to solve moral distributional problems of the sort just mentioned. This is because of the *winner takes all* phenomenon that was discussed earlier in the article.

The second problem is that advances in technology may be such that human designers, programmers and overseers will become less needed over time. Automated technologies may become sufficiently advanced that they can maintain, upgrade and

¹⁸ I defend this view of future governance systems at much greater length in [reference omitted]

develop themselves without human oversight. Thus, even if human creators start these systems, the systems may advance in a way that breaks any meaningful link between those original creators and the outcomes of the system. Indeed, this is already happening, as hinted at by the creative discoveries of the C-Path algorithm. It may also be what happens in the political domain. Humans may still be voted into office to ostensibly oversee technological governance systems, but they may have limited ability to second-guess and intervene.

Finally, even if humans will always participate in such systems, the participation in question has to be of the right type in order to facilitate meaningfulness in the objective or hybridist sense. In current kidney exchange systems, humans still have to make the decision to become donors, and this requires some moral courage. Those who become donors can genuinely claim to make the world a better place. The same is true for the (still largely human) surgical teams that perform the operations. But how long will this be true? If the promise of artificially generating replacement organs using a combination of stem cells and 3D printing comes true, and if surgery itself becomes almost fully automated, the sole remaining participation for human beings will be as *beneficiaries* of automated systems. This may make them *subjectively* better off, but they will no longer be agents or creators of objective outcomes. Again, the necessary causal link is severed.

It is important that this argument is not misunderstood. The claim is that the kinds of machines that make widespread technological unemployment possible are also likely to be better (more accurate, more efficient, less prone to distortion or bias) at attaining the True and the Good. These still remain as sources of objective meaning. And allowing machines to attain those objectives will probably make our lives subjectively better. But the objectives will no longer be *ours*. That is the sense in which the link is severed.

There are two further objections to this argument.¹⁹ The first is that even if machines are better at achieving certain objective outcomes there is nothing to stop humans from achieving them too. If they succeed they will have satisfied the requisite conditions for meaning. This is true, but only up to a point. In the technologically advanced future a human could still develop the ability to, say, design and engineer a

¹⁹ I am indebted to an anonymous reviewer for pushing me on these two issues.

bridge or cure a medical illness, but because the machines are better at doing both of these things there are likely to be a significant personal and social disincentives against doing so: why would you waste time and risk lives if the machines are better?

The second objection is that this argument commits something akin to the Luddite fallacy: it assumes that if machines get better and better at doing things they will take away from the fixed lump of potentially meaningful activities that are open to human beings. But why couldn't more and more objectively meaningful activities open up? This has an air of plausibility but it seems like it is vulnerable to the same objections as the original Luddite fallacy. Some moral and intellectual objectives will suffer from an inelasticity of demand problem: there will be a limit to how much demand there is for insight or resolution, e.g. a grand unified theory of physics or the eradication of hunger and disease. The superstar/winner-takes-all effect will mean that new fields of opportunity will be quickly gobbled up. And exponential progress in technology will mean that humans are not able to redirect their skills and attention fast enough.

In the end, the only domain in which humans might be able to meaningfully contribute to objective outcomes would be in the realm of private, ludic or aesthetic activities, e.g. in producing works of art, or pursuing games, hobbies and sports. In these aesthetic domains, it is less clear that automating technologies help to produce *better* outcomes. We could get computers to create music and visual art or play sports — (in fact, we already can) — but it's not clear that this adds to the aesthetic value in the world or detracts from the value of our own pursuits in those domains. But there is still ample space for automation in our private lives (e.g. in the shape of digital assistants or robot helpers) and these devices could sap us of some of the motivation and ability to engage in acts of private or aesthetic meaning. We also need to ask whether this remaining domain is enough? Metz (2010) argues that the Good, the True and the Beautiful are the three major domains of meaning in human life. In the hypothetical future just sketched, we are left, in effect, with the Beautiful. That certainly looks like a more impoverished form of existence.

5. An Integrationist Approach to Technology

The problem identified at the end of the last section concerns the human relationship with automating technologies. I argued that if meaning and fulfillment are

at least partly determined by how our individual activities join up with the external world, then there is something to worry about if automating technologies takeover most domains of human activity. They could sever the necessary links between what we do and what happens in the world. In making this argument, I was highlighting a potential flaw in our relationship with technology. To a large extent, we still view technology as something external to us: as a tool for shaping and reshaping the world. For most of human history, those technological tools were controlled and manipulated by human hands and minds. So there was a sense in which they empowered our own achievements. But with the growth of automation and AI, this becomes less and less true. The tools have become ever more *externalised*. It is this externalisation that looks like the major threat to continued meaning and fulfillment.

Perhaps there are other ways in which we could relate to technology? Instead of pursuing a strategy of increased externalisation, perhaps we could pursue increased *integration*. In other words, perhaps we could merge ourselves (our bodies; our minds) with technology. Perhaps we could become cyborgs. That way, the systems that I lamented towards the end of the previous section would no longer threaten to sever the link between what we do and what can be achieved. They would be integrated into who we are. We could have the best of both worlds: the benefits of the enhanced capacities of technology along with meaningful participation in the outcomes the technology facilitates. This is admittedly speculative, and I do not have the scope in this article to map out the possible forms of integration, but some of the possibilities are already apparent (Gunkel 2012 & 2015; Haraway 1991; Schermer 2009). I presume the necessary integration would involve increased use of brain-computer interfaces, nanotechnology and various other neuroprosthetic devices. The precise mix and form would need to be debated and specified.

To be clear, I am not claiming here that the integrationist project will be easy and risk-free. Integrating ourselves with technology always poses risks and if the project is to be pursued it must be pursued with due caution and wisdom. What I am claiming is that there could be benefits to be reaped from doing so. That said, there are two more specific critiques of this suggestion that are worth addressing. The first draws its inspiration from defenders of the *extended mind thesis*; the second focuses on a

problematic assumption underlying the entire argument thus far, namely *the 'primacy of the real'*.²⁰

The extended mind thesis is the claim that the human mind already naturally extends into technological artifacts. The concept originates in a paper by Clark and Chalmers (1998). They argued that functionalist theories of mind support the notion that the human mind can be *realised* across multiple physical systems. To support this, they gave a thought experiment involving a man named Otto who had some form amnesia and used a notebook to store and recall various memories. They argued that Otto's memory recall system was functionally equivalent to that of someone who used their brain to store and recall memories. The only difference was that Otto's mind extended into a (primitive and simplistic) technological artifact. Clark and others have developed this idea to suggest that humans are constantly extending their minds, and that devices like smartphones and digital assistants can all form part of their extended minds.

This thesis can be used to challenge, or at least modify, my suggestion that we should pursue an integrationist relationship with technology. When I use the term 'integration', I imagine a future in which technology is directly integrated into biological systems. But according to proponents of the extended mind thesis there is a sense in which we have already integrated ourselves with technology. The AI and automating technologies that I mentioned above are already parts of our extended minds. We don't need to go any further. We have the link we need. The extended mind thesis is certainly an intriguing philosophical claim, but it is dubious as to whether it can restore the link between individual activity and outcomes. For one thing, it is a controversial philosophical thesis with plenty of critics (Menary 2010). More importantly, it is not clear that such technologies give us the right kinds of ownership, control or phenomenological connection with the outcomes they facilitate. There are important phenomenological differences between solving a mathematical puzzle by oneself and doing so with the assistance of a calculator. If we design and programme extended mind technologies for ourselves, then perhaps we could retain some form of the necessary. But as has been argued already, it is not clear whether individuals will do such things in the automated future.

²⁰ I take this term from JP and TK

The second objection is that in endorsing this integrationist approach, and in claiming that technology could sever the link between our activities and the external world, I have been making a dubious axiological assumption. I have been assuming that outcomes in the *real* world are the only type that are sufficient for sustaining a meaningful life. But why must the outcomes be in the real world? One thing we may be able to do in the technological future is create vast virtual worlds in which all manner of activities can be pursued and outcomes realised. Indeed, we could be almost infinitely creative in terms of how we set the parameters for existence in those worlds. Is it not some biased assumption in favour of the real world that denies the meaningfulness of such activities?

This too is an intriguing possibility, but there are some obvious objections to it. Robert Nozick's Experience Machine argument is the best-known (Nozick 1974). The argument is premised on a thought experiment in which you are given the choice between a qualitatively superior virtual-world life and a qualitatively inferior real-world life. Nozick's intuition was that, given the choice, people would choose the latter over the former. The suggestion then being that reality has primacy over virtuality when it comes to what counts for the good life. Although this strikes many as a persuasive argument, recent results in experimental philosophy cast doubt upon the strength of the intuitions Nozick uses to power his argument (de Brigard 2010). Studies suggest that Nozick's thought experiment may simply be exploiting a status quo bias, and that if you change the scenario so that people are asked to 'plug out' of a virtual reality that they enjoy, instead of plugging into one (as in Nozick's original experiment), most people prefer to stick with what they've got.

The full implications of these experiments, and the persuasiveness of Nozick's argument, are still being debated. I find myself drawn to Nozick's original argument, and so I think the integrationist solution remains preferable, but suppose for present purposes I concede that actions in a purely virtual world might suffice for meaning. In that case, it would follow that we need to think carefully about the future into which we are wandering. If virtual reality is our best hope, then this kind of technology needs to be prioritised and made available to all. That may be the only way to solve the deficit in meaning that other forms of technological advancement open up.

6. Conclusion

This article has taken a number of twists and turns. It is probably worth concluding by summarising and drawing together the main lines of argumentation. Four points seem particularly worthy of emphasis. First, assuming that technological unemployment is a real possibility in the not-too-distant future, there are two main ethical and social problems to which it could give rise: (i) the distributional problem, arising from the fact that more income goes to capital owners; and (ii) the problem in personal axiology, arising from the fact that people need to do something with the time they no longer spend working.

Second, I argued although many people celebrate the virtues and pleasures of work and the work ethic, there are reasons to think that working to fill our economic needs is sub-optimal. For one thing, there are a variety of contingent and intrinsic harms associated with work, and for another there are reasons to think that there are better ways in which we could spend our time. This is the antiwork position.

Third, I tempered this enthusiasm for the antiwork position by assessing the possibilities of non-work in light of the four leading theories of what makes for a meaningful and fulfilling life. I argued that there was some reason to be optimistic if you embraced a purely subjectivist theory of meaning; but some reason for pessimism if you embraced an objective or hybridist theory of meaning. The reason for pessimism was that technological advances do not occur in a vacuum: the kinds of technological advance that make structural unemployment possible may sever the connection between human activity and meaningful outcomes across a number of domains.

Fourth, and finally, I suggested that one way to address this problem was to endorse an *integrationist*, as opposed to an *externalist*, approach to our relationship to technology. This would not be a straightforward or easy project, but it may need to be prioritised in order to address the meaning-deficit made possible by advances in automation.

Acknowledgements – I would like to thank Miles Brundage, Jon Perry, Ted Kupper and three anonymous reviewers for feedback on earlier drafts of this paper.

References

- Ackerman, B, Alstott, A., and Van Parijs, P. (2006) *Redesigning Distribution: Basic Income and Stakeholder Grants as Cornerstones for Egalitarian Capitalism*. London: Verso.
- Agar, N. (2015) *The Sceptical Optimist*. Oxford: OUP.
- Autor, D. (2015a). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives* 29(3): 3-30
- Autor, D. (2015b). Polanyi's Paradox and the Shape of Employment Growth'. In *Re-Evaluating Labor Market Dynamics*. Federal Reserve Bank of Kansas City.
- Autor, D. (forthcoming). The Paradox of Abundance: Automation Anxiety Returns. In Rangan, S (ed) *Performance and Progress: Essays on Capitalism, Business and Society*. Oxford: OUP.
- Atkinson, A. (2015). *Inequality: What can be done?* Harvard University Press
- Beck, A. et al (2011). Systematic Analysis of Breast Cancer Morphology Uncovers Stromal Features Associated with Survival. *Science Translational Medicine* 3(108) - 108ra113.
- Black, B. (1986). *The Abolition of Work and Other Essays*. Loompanics Unlimited.
- Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. WW Norton and Co.
- Brynjolfsson, E. and McAfee, A. (2012). *Race Against the Machine*. Lexington, MA: Digital Frontiers Press.
- Campbell, S and Nyholm, S. (2015). Anti-Meaning and Why it Matters. *Journal of the American Philosophical Association*.
- Carr, N. (2015) *The Glass Cage: Where Automation is Taking Us*. The Bodley Head.
- Clark, A. and Chalmers, D. (1998). The Extended Mind. *Analysis* 58:10-23.
- Cowen, T. (2013). *Average is Over: Powering America Beyond the Age of the Great Stagnation*. Plume.
- Crary, J. (2014). *24/7: Late Capitalism and the Ends of Sleep*. London: Verso.
- Csikszentmihalyi, M. (2007). Experience Sampling Method: Measuring the Quality of Everyday Life", Thousand Oaks: Sage Publications Inc., 2007
- Csikszentmihalyi, M. (1997). *Finding Flow: The Psychology of Engagement with Everyday Life*. New York: Basic Books.
- Csikszentmihalyi, M (1990). *Flow: The Psychology of Optimal Experience*. New York: Harper & Row.
- de Brigard, F. (2010). If you like it, does it matter if it's real? *Philosophical Psychology* 23(1): 43-57
- Denning, S. (2015). The "Jobless Future" is a Myth. *Forbes* 4 June 2015 - available at: <http://www.forbes.com/sites/stevedenning/2015/06/04/the-robots-are-not-coming/>

- Dworkin, G. (1981). The Concept of Autonomy. *Grazer Philosophische Studien* 12/13: 203-21
- Fleck, S., Glaser, J. and Sprague, S., (2011). The Compensation-Productivity Gap: A Visual Essay. *Monthly Labour Review*, January 2011, 57 – 69, available at: <http://www.bls.gov/opub/mlr/2011/01/art3full.pdf>
- Ford, M. (2015) *The Rise of the Robots: Technology and the Threat of a Jobless Future*. New York: Basic Books.
- Ford, M. (2009) *The Lights in the Tunnel: Automation, Accelerating Technology and the Economy of the Future*. CreateSpace Independent Publishing.
- Frey, CB and Osborne, MA (2013). The Future of Employment: How Susceptible are Jobs to Computerisation. Oxford Martin School, Working Report.
- Frayne, D. (2015) *The Refusal of Work: The Theory and Practice of Resistance to Work*. London: Zed Books.
- Gheaus, A and Herzog, LM (forthcoming). The Goods of Work (Other than Money). *Journal of Social Philosophy*. Available at <http://philpapers.org/rec/GHETGO> (accessed 13/2/15).
- Gilbert, D. (2005). *Stumbling on Happiness*. Vintage.
- Gilbert, D. T., Pinel, E. C., Wilson, T. D., Blumberg, S. J., & Wheatley, T. (1998). Immune neglect: A source of durability bias in affective forecasting. *Journal of Personality and Social Psychology*, 75, 617-638.
- Gilbert, D. T., Lieberman, M. D., Morewedge, C. K., & Wilson, T. D. (2004). The peculiar longevity of things not so bad. *Psychological Science*, 15, 14-19.
- Gorz, A. (2011). *Critique of Economic Reason*. 2nd Edn London: Verso.
- Graeber, D. (2013). On the Phenomenon of Bullshit Jobs. *Strike! Magazine* - 17 August 2013, available at <http://strikemag.org/bullshit-jobs/>
- Gunkel, D. (2012). *The Machine Question*. MIT Press.
- Gunkel, D. (2015). Resistance is futile: Cyborgs, Humanism and the Borg. In Brode, D. & Brode, S. (eds) *The Star Trek Universe: Franchising the Final Frontier*. New York: Rowman and Littlefield.
- Haraway, D. (1991) *Simians, Cyborgs and Women: The Reinvention of Nature*. New York: Rowman and Littlefield.
- Hughes, J (2014). A Strategic Opening for a Basic Income Guarantee in the Global Crisis Being Created by AI, Robots, Desktop Manufacturing and BioMedicine. *Journal of Evolution and Technology*. 24(1): 45-61.
- International Labour Organisation (2013). Global Wage Report 2012/13: Wages and Equitable Growth, available at http://www.ilo.org/wcmsp5/groups/public/---dgreports/---dcomm/---publ/documents/publication/wcms_194843.pdf
- Kaplan, J. (2015). *Humans Need Not Apply: A Guide to Wealth and Work in the Age of Artificial Intelligence*. Yale University Press.
- Keen, A. (2015). *The Internet is Not the Answer*. Atlantic Monthly Press
- Krugman P. (2012). Robots and Robber Barons. *New York Times*, 9 December.
- Krugman, P. (2013). Sympathy for the Luddites. *New York Times*, 13 June.

- Kurzweil, R. (2006). *The Singularity is Near: When Humans Transcend Biology*. London: Penguin
- Lafargue, P. (1883). *The Right to be Lazy*. Charles Kerr and Co.
- Levine, A. (1995). From fairness to idleness: is there a right not to work? *Economics and Philosophy* 11: 255-274.
- Luper, S. (2014). Life's Meaning. In Luper, S (Eds) *The Cambridge Companion to Life and Death*. Cambridge University Press.
- Maskivker, J. (2010). Employment as a Limitation on Self-Ownership. *Human Rights Review* 12(1): 27-45.
- Menary, R (Ed) (2010). *The Extended Mind*. Cambridge, MA: MIT Press
- Metz, T. (2013). *Meaning in Life*. Oxford: OUP.
- Metz, T. (2010). The good, the true and the beautiful: Toward a unified account of great meaning in life. *Religious Studies* 47(4): 389-409.
- Mokyr, J., Vickers, C. and Ziebarth, N. (2015). The History of Technological Anxiety and the Future of Economic Growth: Is this time different? *Journal of Economic Perspectives* 29(3): 31-50.
- Nozick, R. (1974). *Anarchy, State and Utopia*. New York: Basic Books.
- Packer, G. (2013). Upgrade or Die. *New Yorker*. 4 March.
- Pasquale, F. (2015). *The Black Box Society*. Harvard University Press.
- Pearce, D. (1995). *The Hedonistic Imperative* - available at <http://www.hedweb.com/>
- Piketty, T. (2014). *Capital in the 21st Century*. Harvard University Press.
- Pistono, G. (2012). *Robots will steal your job but that's OK: How to survive the economic collapse and be happy*. CreateSpace Independent Publishing.
- Pratt, G. (2015). Is a Cambrian Explosion Coming for Robotics. *Journal of Economic Perspectives* 29(3): 51-60.
- Rifkin, J. (2014). *The Zero-Marginal Cost Society: The Internet of Things, the Collaborative Commons, and the Eclipse of Capitalism*. St Martin's Griffin.
- Rifkin, J. (1997). *The End of Work*. Jeremy P. Tarcher/Putnam Books.
- Roth, A. (2015). *Who gets what and why?* Houghton Mifflin Harcourt.
- Russell, B. (2004). *In Praise of Idleness*. 2nd Edn, London: Routledge.
- Sachs, J., Benzell, S.G. and LaGarda, G. (2015). Robots: Curse or Blessing? A Basic Framework. NBER Working Paper 21091- April.
- Schermer, M. (2009). The Mind and the Machine: On the Conceptual and Moral Implications of Brain Machine Interaction. *Nanoethics* 3(3): 217-230.
- Schwartz, A. (1982). Meaningful Work. *Ethics*, 92(4): 634-646.
- Smuts, A. (2013). The Good Cause Account of Meaning in Life. *Southern Journal of Philosophy* 51(4): 536-562.

- Sönmez, T. and Unver, U. (2013) Market Design for Kidney Exchange. In Vulkan et al (eds) *The Handbook of Market Design*. Oxford: OUP.
- Srnicek, N and Williams, A (2015). *Inventing the Future: Postcapitalism and a World Without Work*. London: Verso.
- Taylor, R. (2008). The meaning of life. In *The meaning of life*, ed. E.D. Klemke and Steven M. Cahn. New York: Oxford University Press
- van Parijs, P. (1995). *Real Freedom for All: What (if anything) can justify capitalism?* (Oxford: Clarendon Press.
- Vulkan, N., Neeman, Z and Roth, A. (Eds) (2013). *The Handbook of Market Design*. Oxford: Oxford University Press
- Weeks, K. (2011). *The Problem with Work: Feminism, Marxism, Antiwork Politics and Postwork Imaginaries*. Duke University Press.
- Wertheimer, A. (1988). *Coercion*. Princeton NJ: Princeton University Press.
- Widerquist, K. (2013). *Independence, Propertylessness and Basic Income: A Theory of Freedom as the Right to Say No*. New York: Palgrave MacMillan.
- Widerquist et al (eds) (2013). *Basic Income: An Anthology of Contemporary Research*. Oxford: Wiley-Blackwell.
- Wielenberg, E. (2005). *Value and Virtue in a Godless Universe*. Cambridge University Press.
- Wolf, Susan. 2010. *Meaning in life and Why it matters*. Princeton, NJ: Princeton University Press
- Woodcock, G. (1944). The Tyranny of the Clock. *War Commentary for Anarchism*.
- Zarsky, T. (2012). Automated Predictions: Perception, Law and Policy. *Communications of the ACM* 15(9): 33-35
- Zeckhauser, Robert (1973). Time as the Ultimate Source of Utility. *Quarterly Journal of Economics* 87: 668-673.