Technological Change and Human Obsolescence

An Axiological Analysis

By John Danaher, NUI Galway

Abstract: Can human life have value in a world in which humans are rendered obsolete by technological advances? This article answers this question by developing an extended analysis of the axiological impact of human obsolescence. In doing so, it makes four main arguments. First, it argues that human obsolescence is a complex phenomenon that can take on at least four distinct forms. Second, it argues that one of these forms of obsolescence ('actual-general' obsolescence) is not a coherent concept and hence not a plausible threat to human well-being. Third, it argues that existing fears of technologically-induced human obsolescence are less compelling than they first appear. Fourth, it argues that there are two reasons for embracing a world of widespread, technologically-induced human obsolescence.

Keywords: axiology; technology; human obsolescence; value theory; virtue; the good life; enhancement; robotics; artificial intelligence

1. Introduction

Imagine waking up one morning and finding out that you are no longer useful. The skills that you spent years developing, the capacities you finely honed through hours of practice, are all yesterday's news. Thus consigned to the trash-heap of history, how would you feel? Sad? Angry? Relieved?

It may sound fanciful, but with rapid technological progress answering these questions has become a pressing concern. As new technologies develop,

old ones obsolesce. Since our success in life depends largely on how skillfully we wield these technologies we run the risk of becoming obsolete alongside them (Sparrow 2015 & 2019; Currier 2015). Indeed, this has already happened to some of us. Generations of skilled workers, for example, have seen their economic usefulness ebb away in the face of rampant workplace automation: the skills they developed no longer suit the updated technological environment (Frey 2019). As we turn the lens of technological development on ourselves, and see our bodies and brains as among the technologies that we can tinker with and improve, this tendency toward obsolescence is likely to grow. What impact will this have on the value of our lives?

This question has received some recent attention in the philosophical literature (Sparrow 2015 & 2019; Danaher 2019b), but that attention has been largely negative. The common presumption seems to be that technologically-induced obsolescence is a bad thing and we ought to do what we can to mitigate its worst effects. In this article, I offer a more positive outlook. I argue that technologically-induced human obsolescence can be welcomed *if* we can build a culture that embraces its advantages.

I defend this view in four stages. First, I clarify what is meant by technologically-induced obsolescence, identify the primary mechanisms of obsolescence, and classify its four basic forms. Second, I consider the coherence of the claim that humans can be rendered obsolete by technology, arguing that there are versions of this claim that are plausible and others that are not. Third, I examine and challenge the claim that technologically-induced obsolescence is a bad thing, focusing in particular on arguments presented by Robert Sparrow against enhancement-induced obsolescence. Fourth, and finally, I present two arguments for thinking that technologically-induced obsolescence is something to be welcomed.

2. Understanding Human Obsolescence

Let's start with the nature of obsolescence. I will discuss this initially from a value-neutral perspective. In other words, I will try to clarify the nature of the phenomenon and not pass judgment on whether it is a good or bad thing. Adopting this stance of value-neutrality is tricky since the word 'obsolescence' carries pejorative connotations for most people. Nevertheless, it is essential that this stance of value-neutrality be accepted if the analysis is to be properly appreciated. For example, I will, in this section, be mentioning disability and old age as examples of obsolescence. This will strike many people as an offensive thing to say. Old people and disabled people are not obsolescent! I quite agree. People who are inclined towards this interpretation of my comments should bear in mind the stance of value-neutrality that is being adopted. They should also bear in mind that later in this article I will be defending the idea that building a culture of obsolescence is a good thing and that we should reject the notion that the value and purpose of human life is assessed in terms of whether it is obsolescent or not. In this respect, my analysis of the relationship between technology and obsolescence might be seen to be consistent with recent critical analyses of the relationship between technology and disability, specifically those that urge us to move away from the belief that disabled people can be liberated or redeemed by technology (e.g. Nelson et al 2019; Shew 2019; Shew 2018).

With that caveat in mind I can proceed to analyse the concept of obsolescence. Obsolescence is the state of no longer being useful or used. It is a relative phenomenon. To determine whether something is obsolescent there must be some set of goals/standards against which its usefulness can be measured. In measuring this usefulness, we can distinguish between obsolescence as a *process* and obsolescence as a *state of affairs*. Typically, something will go through a *process of obsolescence* and reach *the state of affairs of being obsolete*. This will happen if the thing was once able to meet the relevant goals or standards, but no longer can, or, alternatively, if it is no longer the most efficient or effective means of achieving those goals or meeting those standards. That said, it may also be possible for something to be obsolete simpliciter, without ever having gone through a process of obsolescence. This would happen

if the thing was never able to reach the relevant standards or achieve the relevant goals.

Whether it be a process or a state of affairs, obsolescence can never be determined solely by examining the intrinsic functionality of an object or device but, rather, only by looking at its comparative functionality. For example, your thirty-year old computer may still function adequately in some respects -- e.g. you may be able to use its word processor or calculator -- but will undoubtedly face problems in other respects -- e.g. it may no longer be updated with the latest software, it may have become vulnerable to many computer viruses, and it may not connect to the internet and communicate with other devices. More recent models will be much more effective and efficient means for achieving your goals because they can smoothly integrate with the contemporary technological ecosystem. Your model has consequently been rendered obsolete by changes in this broader technological ecosystem and not due to its inherent features.

How does this relativised understanding of obsolescence apply to human beings? To put it briefly, human obsolescence arises from a *mismatch* between the *internal capacities* of the human being, and the *external standards/goals* of the society surrounding that human being. Each human being has some set of cognitive, emotional and physical capacities that enables them to perform to a certain standard and achieve certain goals. These standards and goals are set by economic, political, social and personal forces.[1] The capacities are partly innate and partly a function of how the individual is trained and educated. Obsolescence arises when the internal capacities are no longer able to meet the desired standards or achieve the desired goals. To give an example, humans who are really good at performing arithmetical functions, either in their heads or with the

[1] Here is another place where my discussion of obsolescence links up, to some extent, with debates about disability. My definition of obsolescence, as something that arises from a mismatch between intrinsic capacities and external standards is somewhat similar to the claim made by the various 'social models' of disability (i.e. models that claim that disability is not something intrinsic to the person but rather something that is, at least in part, socially determined).
[2] Congenital disability might be a different case since there is no decline in capacity over a lifetime. The mismatch between intrinsic ability and external standards may exist from birth. It is important to bear in mind, however, that

aid of pen and paper, were once highly demanded in scientific and military research (so-called 'human computers'). They have since been rendered obsolete by advances in digital computing technology. This is not because these humans lost their internal capacity for arithmetic; it is rather because computing technology raised the expected performance benchmarks to a level that humans could no longer match. As a result, human computers were no longer deemed economically or socially useful, and there was no incentive to train the skill into future generations. This is a classic example of technologically-induced human obsolescence.

Not all forms of obsolescence are the same, and not all are technologically-induced. There are at least three primary mechanisms of human obsolescence:

- Obsolescence can occur if there is some decline in the individual's internal capacities (relative to their own historical personal baseline), such that they are no longer able to achieve the relevant goals or meet the preferred standard. This is something that can happen through accident, illness or old age. It is also, controversially, something that can happen through acquired disability.[2] The examples of old age and disability are particularly important and will be considered in more detail later in this article.
- Obsolescence can occur if there is some change in the external standards and goals. In other words, economic markets and social institutions can change in such a way that the internal capacities of the individual no longer match the external needs. This can be due to technological change (e.g. as in the case of the switch from agricultural to industrial societies from the 1750s onwards) or other cultural and social change. For

---

[2] Congenital disability might be a different case since there is no decline in capacity over a lifetime. The mismatch between intrinsic ability and external standards may exist from birth. It is important to bear in mind, however, that obsolescence is not necessarily the same thing as disability. Someone could count as disabled without being obsolescent (according to my definition) and vice versa. If the social model of disability is adopted, there may be more overlap between the concepts.

example, a Catholic priest may find their social usefulness under siege due to the waning influence of the Church in developed societies.

- Obsolescence can occur if there is some improvement in the capacity of rivals (i.e. other humans or machines) such that the individual's performances, although still matching their personal baseline, are comparatively worse. This is what happened to the human computers when digital computers came along and it is a particular problem in the case of competitive goods or practices.

Particular episodes of obsolescence may, of course, result from a combination of the three mechanisms.

There are three key variables that moderate the impact that obsolescence has on the individual. Two of these have previously been described by Robert Sparrow (2015 & 2019). The first is the *pace of change*: in other words, how fast rivals are getting better, standards changing, and capacities declining. The faster the pace of change, the more disruptive it is likely to be to the individual's sense of worth and social meaning; the slower the pace of change, the more likely it is that the individual can adjust to the new reality. Alvin and Heidi Toffler's 1970 book *Future Shock* is a classic exploration of this problem, describing in some detail the existential angst that occurs when humans cannot keep up with the pace of social change.

The second variable is the *adaptability and updateability* of the human being: in other words, whether they can retrain or reskill themselves in such a way that they can offset the impact of obsolescence or, at the very least, find some other calling in life in which they are still capable of performing to the desired standard. If they can, then the disruptive impact of obsolescence can be reduced.

The third variable is the adapatability of the society around the individual: in other words can society recognise the individual's obsolescence in a particular domain and then adapt and find other outlets for that person's capacities. When artificial computers rendered human computers obsolete, did

society find other job opportunities and outlets for these individuals? Could it change its norms and practices to accommodate these newly obsolete workers? The adaptability of society, as we shall see below, might be the key to mitigating the worst effects of obsolescence.

Not only are there different mechanisms underlying obsolescence, there are also different forms that the resulting obsolescence can take. Four are particularly important when it comes to assessing the axiological impact of obsolescence. They are:

- **Actual Obsolescence**: This arises when an individual experiences some genuine and objectively determinable form of obsolescence, i.e. their internal capacities can be shown to no longer meet the desired performance standards or achieve the desired goals.
- **Perceived Obsolescence**: This arises when the individual believes themselves (or is believed by others) to be obsolete. This can occur in the absence of actual obsolescence. For example, elderly people and disabled people are are sometimes falsely assumed to be obsolescent with respect to certain economic skills when they are actually perfectly capable of performing to the relevant standard. They experience perceived obsolescence but not actual obsolescence.
- **Narrow Obsolescence**: This arises when the individual undergoes obsolescence with respect to some specific or narrow range of activities and skills.
- **General Obsolescence**: This arises when the individual undergoes obsolescence with respect to a broad range (conceivably all known) activities and skills.

The line separating narrow and general obsolescence is not a sharp one; it is blurry and the two fade into one another. In a sense then they form a single dimension along which cases of obsolescence can vary. Obviously, the most extreme form of general obsolescence arises when the individual is no longer capable of anything, but presumably there could be cases in which someone is capable of meeting one or two external performance standards but otherwise

suffers from such widescale loss in ability that it might be right to refer to their case as one of general or at least extremely broad obsolescence. The actual-perceived distinction is somewhat different insofar as it focuses on the difference between subjective beliefs and objective reality. That said, what counts as the objective reality is often determined by intersubjective agreement so the line between the two may also be somewhat blurry and they may form a single dimension of variance. In the next section we will arrange these four types of obsolescence into a simply 2-by-2 matrix that supports this 'dimensional' interpretation.

Each of these four forms of obsolescence has different potential axiological impacts. Prima facie, the worst kind of obsolescence would be *actual-general* obsolescence because it implies some catastrophic and genuine loss in the ability to participate in society. Narrow and perceived forms of obsolescence look less concerning because they often have readily-identifiable solutions: the human can switch to another domain of activity or, in the latter case, try to change their perception of the situation. But we shouldn't be too hasty in reaching such conclusions. Changing perceptions and changing domains of activity can be practically difficult. Furthermore, one of the major goals of this article is to argue that even very broad forms of obsolescence might be better than we are initially inclined to think.

Before we reach this conclusion, however, we first need to conduct a deeper investigation into the coherence of obsolescence. Is it really possible for humans to obsolesce in a general or narrow sense?

3. Is Human Obsolescence Possible?

It might be argued that it isn't coherent to suggest that humans can obsolesce. Humans, it could be said, do not serve purposes and so to suggest that they could be rendered obsolescent by technology is both implausible and non-sensical. Is that right? Answering this question is important. If technologically-induced obsolescence is not a plausible or coherent notion, then there is little

point in addressing its broader axiological impact. If it is, then this investigation is worthwhile.

To assess the plausibility and coherence of human obsolescence, we need to consider the different forms of obsolescence outlined in the previous section. Doing so, we see that there isn't just one question to deal with here; there are four. We are not just asking whether obsolescence, in general, is coherent and plausible; we are asking whether actual-general, actual-narrow, perceived-general and perceived-narrow forms of obsolescence are coherent and plausible. For ease of analysis, we can arrange these four possibilities into a two-by-two matrix. I have done this in Table 1 and suggested some answers to the question of whether each is plausible and coherent.

	Actual	**Perceived**
General	Uncertain	Plausible and coherent
Narrow	Plausible; historically confirmed	Plausible; historically confirmed

Table 1 – Forms of Obsolescence

As you can see, I believe that all but one of these four forms of obsolescence is plausible and coherent. The one exception is actual-general obsolescence, which I have earmarked as being of uncertain plausibility. Let me explain the reasons for my assessments of the four possibilities in what follows.

I'll start by explaining why I think both of the narrow forms of obsolescence are plausible. The obvious reason is that both have already happened. The examples given in the previous section illustrate this. For instance, one could argue that the history of workplace automation reads as a litany of narrow forms of technologically-induced obsolescence (Danaher 2019a; Frey 2019). Some of these are actual in nature: specific human capacities are no longer objectively useful in the workplace; some of these are perceived: the

9

human-performed skill still has some utility but is perceived to be useless by the individual and the society around him/her and so is not developed any further.

The two forms of general obsolescence are the trickier cases. To suggest that humans could become generally obsolescent commits the fallacy of supposing that humans serve a finite number of purposes and so as soon as they lose the capacity to serve all those purposes (either due to internal or external factors) they become generally obsolescent. This seems like it must be false. Humans do not serve a finite number of purposes. Humans are an adaptable and flexible species. Humans find entertainment and a sense of meaning in a diverse range of activities. Humans are constantly finding new activities and new purposes. There doesn't seem to be any obvious limit to the purposes they can serve. Furthermore, even when human capacities do suffer from some comparative or absolute decline relative to, say, machines, they still find it meaningful to engage in certain activities. Game-playing is a classic example of this. It has long been the case that machines have been capable of outperforming humans in chess. **In the strict sense, then, humans are actually obsolete when it comes to chess playing**. But human chess-playing tournaments are still very popular and competitive. People work hard at improving their abilities in chess and derive a lot of satisfaction and meaning from doing so. Consequently, at least in this one domain, technological inferiority hasn't resulted in a loss of human opportunity or of an outlet for human flourishing.

It is worth reflecting in more detail on why this might be true. Remember, it is not because humans can never become obsolete in certain activities -- the historical examples of narrow obsolescence prove that they can -- it is because some human activities, such as chess, are not defined or valued in terms of the instrumental value they serve. Compare games with workplace activities. Workplace activities are primarily valued for their instrumental purposes: do they make a company more profitable, do they contribute to economic growth? If machines can achieve those ends more efficiently and effectively than humans, or if humans decline in their capacity (relatively or absolutely) such that they no longer meet the performance benchmarks, then obsolescence in the workplace is possible. This does not mean that workplace activities have no intrinsic purpose

or value; it just means that if they didn't serve the instrumental (and economic) needs of the individuals and corporations in question, they would not required. But games are different. They are not valued primarily for the instrumental ends they serve (e.g. the fame, glory or wealth of the participants) but for the kinds of skilled performances that are intrinsic to the game and the pleasure they give to the performers. As long as humans are capable of participating in those performances, they will not be rendered obsolete from those domains of activity. The idea of obsolescence simply makes no sense when applied to them. Presumably, something similar will be true for other non-game domains, such as nurturing friendships and family relationships. Although they may serve instrumental ends, they are not primarily valued in terms of those instrumental ends.

The bottom line, then, is that general forms of obsolescence seem to be implausible because (a) there doesn't appear to be a finite number of activities and purposes that humans serve (or, at the very least, the set of activities is larger than humans can practically exhaust); and (b) at least some human activities are resistant to the dynamics of obsolescence because they are valued for their intrinsic performances and not their ultimate ends.

But why then did I say that actual-general obsolescence is of uncertain status as opposed to saying that it is genuinely implausible? There are three reasons for this. First, in some extreme cases, actual-general obsolescence is a possibility. For example, certain individuals may experience a catastrophic loss of capacity due to accident or severe disability (e.g. someone in a persistent vegetative state). This individual may be generally obsolescent, for all intents and purposes. That said, these cases are likely to be rare and not attributable to technologically-induced changes.[3] Second, even though it seems plausible to suggest that we can always find new purposes and new non-instrumental activities to while away the hours, this may be more true in principle than in practice. Society might struggle to adapt to rapid technological change and create

---

[3] It is also true to say that in the extreme case we all die and thus become generally obsolescent, but presumably that shouldn't count as an example of obsolescence since one must still exist in order to count as obsolescent.

new opportunities. For example, at the moment we live in societies that are deeply wedded to the work ethic and the productive demands of capitalism. This is a society in which fears of obsolescence are palpable. There is a significant amount of cultural baggage and institutional inertia built up around these practices. If there is some rapid, technologically-induced obsolescence with respect to the human capacity to contribute to the capitalist demands, these institutions may not be able to adapt to the new reality. Consequently, for all practical purposes, some people may be rendered generally obsolescent within our current social order (in Marxist theory this is referred to as the problem of 'surplus populations', i.e. populations of people not needed for capitalist production – Srnicek and Williams 2015). Third, as will be discussed in more detail below, the problem of obsolescence may manifest itself even in non-instrumentalist domains of activity, if individuals suffer from comparative loss relative to the other humans (not machines) participating in those activities.

It is worth bearing in mind that the scepticism about general obsolescence applies only to actual-general cases of obsolescence. General-perceived obsolescence is a much more tangible possibility. Even if there are other things humans could do to find a sense of purpose or meaning, they might not be able to see this or they may be actively encouraged not to see it. The dynamics of obsolescence in old age are an obvious case in point. Several philosophers have pointed out that there is a tendency in Western societies to assume that the elderly have nothing to contribute to society, even when this is not actually true (De Beauvoir 1996; Baars 2012 & 2017; Harriott 2006). The result of this widespread assumption is that both they and their family and friends perceive them to be obsolescent: to be a burden on the younger population. There is pushback against this perception, and there are various attempts to have the capacities of the elderly valued and recognised, but the perception is still out there. Similar problems of perception affect disabled persons.

In conclusion, then, at least three of the four forms of obsolescence are plausible and coherent. Furthermore, the fourth, while perhaps not plausible in principle, may become a practical reality due to institutional and cultural inertia.

Consequently, an investigation into the axiology of human obsolescence in its general and narrow forms is justifiable. In the remainder of this article, I will conduct such an investigation, focusing specifically on obsolescence that is induced by technological change.

4. The Case Against Technologically-Induced Obsolescence

Extended discussions of the axiological impact of technologically-induced human obsolescence are rare. This is not to say that they are non-existent, or that one cannot infer opinions about obsolescence from sources that do not discuss the phenomenon directly.

For example, there are several books and articles that discuss the impact of technology on specific human values (e.g. the impact of AI on equality or the impact of human enhancement on autonomy), and one can generalise from these to reach conclusions about obsolescence. For instance, people who worry about the impact of AI on income inequality are often, implicitly, worrying that disadvantaged groups will be rendered economically obsolete as a result and that this will be a bad thing (e.g. Brynjolfsson and McAfee 2014; Ford 2015; Danaher 2019a).

Direct discussions of the axiological impact of technologically-induced obsolescence are much more limited. One potential exception is Agar (2013 and 2019) who looks at the impact of both enhancement technology and digital technology on human well-being. He worries, in particular, that certain forms of technological and digital enhancement of ourselves and our economy will leave humanity as we know and value it behind. In response to this, he argues that we need to either resist some radical technological transformations (Agar 2013) and/or find a complementary space for humans alongside technology in order to maintain our flourishing and well-being (Agar 2019). Another potential exception is Danaher (2019a and 2019b), who discusses the impact of robotics and AI on the human capacity for agency and who also explicitly begins one of his discussions with the observation that 'human obsolescence is imminent'

(2019a, 1). The gist of Danaher's argument is that technological advances in automating technologies tend to obviate the need for human action and initiative in the world, and thus reduce humans to moral patients: passive recipients of the benefits that technology can bring, not active shapers of the world. If this trend continues to an extreme, this could result in a widescale form of obsolescence. Danaher argues that this would have a devastating impact on the current (Western-liberal) value system since that value system currently prioritises and valorises human agency and responsibility. That said, Danaher never fully fleshes out what he means by obsolescence nor does he continue the theme of obsolescence throughout his work. Instead, to the extent that he can be said to critique technologically-induced obsolescence, he focuses on specific threats to specific human values (such as autonomy, agency and responsibility), and not the phenomenon of obsolescence itself.

These examples notwithstanding it is Robert Sparrow who has developed the most extended and concrete analysis of technologically-induced obsolescence to date (2015 and 2019). His analysis focuses specifically on a form of obsolescence that is induced by advances in human enhancement technology. He claims that rapid improvements in the technology of human enhancement will result in a comparative decline in the capacity of individual humans to reach external performance benchmarks. In other words, if a particular human subject (S1) chooses to undergo enhancement at a particular moment in time (T1), and if the technology of human enhancement rapidly improves over time, then S1 will find himself at a comparative disadvantage relative to S2 who chose to enhance at T2 (who will, in turn, find himself at a comparative disadvantage to S3 who chose to enhance at T3). If the technological improvements come quickly, and if individuals are unable to update their enhancements in line with these changes, they will soon obsolesce. This will have a range of devastating axiological impacts, five of which feature prominently in Sparrow's writings.

First, Sparrow argues that rapid human obsolescence will exacerbate the worst features of capitalism. In capitalist societies people are largely valued for their economic productivity: if they cannot contribute to the economic bottom line, they are no longer seen to be socially useful. Rapid human obsolescence will

result in increased competition for productive jobs and opportunities, increased perception of uselessness, compressed periods of peak productivity, increased intergenerational inequity (because later generations of enhanced humans have an advantage over earlier ones), an overall loss of individual resiliency, and increased dependency on others.

Second, Sparrow argues that the dynamics of rapid obsolescence will take a toll on the individual's subjective outlook. There will be an increased sense of regret associated with the decision to undergo enhancement. No matter what point in time you pick to enhance yourself (or no matter what time others decide for you), you will always have the sense that you enhanced too soon, that you should have waited a little longer until the next round of technological improvements came onboard. This will give rise to a more serious variation on the "option regret" people sometimes experience when they buy a smartphone model that is quickly superseded by the latest generation of devices. Sparrow argues that this option regret will undermine our sense of gratitude -- we will feel more responsible for our own failures and less grateful for what we have.[4]. It will also result in increased status anxiety, as we become more conscious of the precarious nature of our social standing: soon to be nullified by the latest generation of enhanced humans.

Third, Sparrow argues that this type of obsolescence is a particular problem for certain forms of human enhancement. For example, human life extension and genetic enhancement. It is much worse, according to Sparrow, to live a greatly extended life in a state of obsolescence than to live a normal one. Imagine living to be 1,000 years old but spending 950 of those years being obsolescent in the eyes of the society in which you live. It doesn't sound very nice, does it? Similarly, Sparrow argues that genetic enhancements for our offspring will be a particular problem if those enhancements rapidly obsolesce. If children quickly become obsolete, they will grow up resentful of the fact that they were born too soon and will acquire the sense that they are all 'yesterday's children' (Sparrow 2019).

[4] This, of course, is a common objection to the project of human enhancement. See, for example, Habermas 2003, Sandel 2007 and Nagel 2010.

Fourth, Sparrow argues that obsolescence poses axiological challenges even when enhancements are pursued for absolute or intrinsic reasons. Some proponents of enhancement argue that there are intrinsic reasons to favour enhancement -- that it is desirable to be enhanced in and of itself, not just because of the positional advantage that the enhanced have over the unenhanced. They claim that it is a good thing to have enhanced cognitive capacity, or an enhanced sense of well-being, even if your enhancements don't bring you up to the level of some other super-enhanced individual. For example, we can derive intrinsic satisfaction from our improved ability to solve Sudoku puzzles; we don't need to win the Sudoku world championships to have this satisfaction (this is somewhat analogous to the chess example given earlier). But Sparrow argues that it is not so easy to embrace the intrinsic virtues of enhancement. Even if certain enhancements are intrinsically valuable, it will be hard to avoid being aware of the fact that you are comparatively worse than others with respect to those enhancements. Even if you can yourself solve the Sudoko puzzle you do so in the shadow of those who can solve them faster. In this regard, Sparrow's argument is similar to one made Robert Nozick regarding the nature of self-esteem (Nozick 1974; Mason 1990). Nozick argued that our sense of self-esteem is almost always derived from a comparison of our abilities with those of others: if we rank highly relative to others we will have a high sense of self-esteem; if we rank lowly we won't. Sparrow is suggesting that something similar will be true with respect to our intrinsic satisfaction with our enhancements.

Fifth, and finally, Sparrow argues that it is difficult to avoid the problems of obsolescence if you wish to pursue the goal of human enhancement. This is because the desire to technologically enhance ourselves is driven, in part, by the belief that we (in our current form) ought to become obsolescent. It is not easy to sidestep this by arguing that certain forms of enhancement could be used to mitigate the negative consequences of obsolescence (e.g. moral enhancement that might make us more egalitarian and enlightened about the fate of our obsolescent peers), or by hoping that there will be some general social reform that minimises the impact of obsolescence. These hopes are speculative and

contrary to the underlying ethos of enhancement which is to improve that which is imperfect.

These five reasons lead Sparrow to conclude that enhancement-induced obsolescence would be a very bad thing for humanity. He suggests that we should avoid the narrative of "naive inevitabilism" about the march of technological progress, and instead consider taking regulatory steps to prevent or slow down the march toward human obsolescence (Sparrow 2015).

Although Sparrow traces out the consequences of enhancement-induced obsolescence in some detail, his analysis is limited in crucial respects. He focuses on just one potential mechanism of obsolescence to the exclusion of others. He imagines a world in which humans themselves are the technology that is being rapidly improved and so it is relative to other human beings that we are rendered obsolescent. With an eye to human history, this is not the most common mechanism of technologically-induced obsolescence. It is not usually relative to one another that we become obsolescent (except in the case of ageing). It is, rather, relative to our technologies, machines and social orders that we become obsolescent. Human biology and physiology has remained remarkably consistent over the past 40,000 years or so. Despite all the stürm and drang from philosophers and bioethicists about the prospect of radically enhanced human beings, it is primarily through external technological improvements that we 'enhance' ourselves. This is not to say that human enhancement is impossible or will never happen -- the development of CRISPR Cass 9 and its implications for genetic engineering are tantalising -- but it is to suggest that if we are worried about the prospects of widespread human obsolescence, we may be looking in the wrong place if we focus our attention on human enhancement.

Rapid advances in automating technologies (i.e. robotics and AI) combined with associated cultural changes in how we organise and run our societies to take advantage of these technologies, seem to pose a much more credible threat. Consider an example. If autonomous vehicles have the safety benefits that are claimed by the manufacturers, and if they are cheaper and more

efficient than human-driven vehicles (big caveats), then we can expect human drivers to become relatively obsolete and we can expect a change in cultural and legal norms associated with human driving. In the extreme, human-driving may start to be seen as ethically on a par with drink-driving: something that just shouldn't happen in a well-functioning society. Less extreme, human-driving may be heavily discouraged, like speeding is nowadays. Humans may start to see their own participation in driving as a moral luxury and chastise themselves and one another for indulging in it (Nyholm and Smids 2018; Nyholm 2020). We can also expect typical economic logic to take over in the field of driving-related employment. Employers will be incentivised to drop costly and risk-prone human drivers in favour of their robotic equivalents. The net result is that humans will be rendered obsolete (for actual and perceived reasons) from this domain of activity. Multiply this over and over again for all the possible use cases for advances in automating technologies, and you have a credible and looming threat of widescale human obsolescence that has nothing to do with human enhancement.

This matters because the mechanism of obsolescence can be expected to make some difference to the axiological impact of that obsolescence. Humans are naturally social and competitive (Erdi 2019). We instinctively compare and rank ourselves relative to one another. To be rendered obsolete by our fellow human beings has deeper, instinctual resonance than does being rendered obsolete by machines. Enhanced obsolecence means we have lost ranking relative to one another, and not relative to some external force. There might be a better chance of building a narrative of community solidarity around machine-induced obsolescence than there is around enhancement-induced obsolescence.

Consider an example. Imagine you have a factory full of human workers. Suppose the factory goes out of business because the product they make is now obsolete. Would that be worse for the community of workers than if, say, half of them had been fired for perceived under-performance? While both scenarios are tragic for the workers, we might expect the latter to fragment the community and create bitter competitions and resentments. The former, on the contrary, might actually foster greater solidarity: since all workers have been equally

affected by the external technological changes there are legitimate grounds for favouring group empathy and solidarity to get through this difficult time. A similar dynamic might be expected to arise in the case of a group of drivers rendered obsolete by a fleet of autonomous vehicles -- more so than if a subset are rendered obsolete relative to enhanced human drivers anyway.

The point here is twofold: (i) Sparrow's critique of obsolescence focuses on the relatively speculative threat of enhancement-induced obsolescence, and not on the more credible threat of automation-induced obsolescence; and (ii) some of the negative axiological impacts he identifies, specifically those having to do with lost or constrained individual productivity, excess competition and status anxiety, might be less warranted in the latter case. This doesn't mean that they won't arise as a matter of fact; it just means that there is less normative justification for them. You are not being rendered obsolete because of some individual failing and senescence. You are being rendered obsolete for reasons beyond your control. This may be tragic but because the obsolescence comes from an outside and non-human force, there is more reason to adapt and change your expectations as to what it is that gives value to your life. This supports the view that the mechanism of obsolescence makes a difference to its axiological impact. I elaborate more on what it might mean to change your expectations regarding value in your life in the next section.

There are a two other critical observations to be made about Sparrow's argument. First, Sparrow does not specify which form of obsolescence he imagines will take hold in his hypothetical future. Is it actual-general or actual specific, or some other variant? It is unlikely to be actual-general obsolescence. Even in the case of enhanced but obsolete humans, there are still plenty of things they can do and be useful for. They can still form friendships and intimate relationships with one another; they can still play games with one another; they can still drink and be merry. Perceived-general obsolescence may be more likely, particularly if the society in which these enhanced humans live remains wedded to competitive and positional ideologies, such as capitalism. The narrow forms of obsolescence are, of course, even more plausible, but if they are very narrow then the axiological impact will be minimal: the negative consequences to which

Sparrow points will only arise in certain domains of activity but then all the enhanced humans have to do to avoid them is switch to another domain.

We can suppose that Sparrow is not limiting his argument to a very narrow form of obsolescence -- his argument wouldn't carry much weight if it did. So he must be imagining a broader form of obsolescence -- obsolescence across many, many domains of life -- perhaps bordering on perceived-general obsolescence. But since that form of obsolescence is largely a function of the prevailing social ideologies, then, as I will argue below, the key to addressing his worries may lie not so much in reversing the tide of technological progress than it does in reforming those social ideologies.

This brings me to a final point I wish to make about Sparrow's argument. While Sparrow is correct in saying that we should avoid 'naive inevitablism' about technologically-induced obsolescence, it is equally important to recognise that we do not have as much control over the course of technological development as we would like. Sparrow's own view is that the problem of enhanced-obsolescence arises from a collective action problem: the decision to enhance can seem individually rational and beneficial, but if everyone does it we end up in the enhanced 'rat race' that has all the negative features described above. This is why Sparrow calls for regulatory intervention: we need it to correct for the flaws in individual rationality.

Sparrow is probably right about this: collective action problems probably do lie at the heart of this issue. But they go deeper than Sparrow supposes, particularly when it comes to automation-induced obsolescence. It is not just at the individual level that the collective action problem arises: it is also at the institutional and societal level. Particular companies and organisations can rationally favour automating technologies because they give them some advantage over rivals. The same is true for particular communities and governments (partly because automating technologies are associated with economic growth and innovation). Top-down regulatory interference consequently isn't always a credible solution to the problem automation-induced obsolescence since the top-down regulators are implicated in it. Of course, this

doesn't mean that the position I defend in the next section is immune from collective action problems. It suffers from them too. The point is that it is hard to address the societal impact of technologically-induced obsolescence. Given this, perhaps should try to direct our energies at the 'solution' to the problem that is most socially desirable, not the one that is most conservative or obvious?[5]

I'll now argue that the most desirable solution might be to embrace obsolescence, not to halt it. This will call for a radical reconceptualisation of obsolescence. Instead of viewing it as something that is inherently negative or pejorative there is an opportunity to reframe it as a positive. Or so, at least, I shall argue.[6]

5. The Case for Technologically-Induced Obsolescence

There are two arguments for welcoming technologically-induced obsolescence. The first argument invites the reader to see the potential upside to obsolescence; the second argues that building a culture of obsolescence could be a desirable response to the problem of rapid technological advance. In both cases, the arguments speak to the advantages of broad forms of obsolescence (both actual and perceived) and not just narrow forms. It is relatively easy to make the case for the desirability of narrow forms of obsolescence: just point to unpleasant activities that humans are currently under pressure to perform and

[5] Cf. Agar 2013 and 2019 whose solution to the problem of humanity being 'left behind' focuses more on preserving and conserving our current set of values over and above radical techno-social transformation. For example, in Agar 2019, he favours developing a 'social economy' that provides a continued place for humans alongside a 'digital economy' that doesn't. This avoids the more radical economic upheaval that could be caused by embracing economic automation in a more thoroughgoing way.
[6] There are other objections one could make to Sparrow's argument. Many of these take particular umbrage at how he interprets and understands the enhancement project. For example, his suggestion that the enhancement project is necessarily committed to the dynamics of competition could be challenged. Enhancement is a fuzzy concept but a reasonable case can be made that some forms of enhancement may not require an improvement in capacities relative to others and so do not require competition and positionality. For example, enhancement that focuses specifically on subjective well-being and contentment.

argue that it would be good if we could be rendered obsolete from those activities; it is more challenging to make the case for broader forms of obsolescence. That said, both arguments should be read with the earlier caveat in mind: complete actual-general obsolescence is probably not a fully coherent notion because of the diversity of ends that humans can pursue and the fact that some activities are valued for their intrinsic performances and not for their ends. Indeed, both of these facts will feature in the arguments given below.

The first argument is simple. Recall that obsolescence arises when there is a mismatch between internal capacities and external demands/standards. Criticisms of obsolescence, such as those made by Sparrow, dwell on the fact that these external demands and standards are an important source of value. Consider the fact that many people find meaning by making contributions to their societies through their activities (be they economic, political, charitable or otherwise in nature). It is only if they can make those contributions that they (and others) will attach value to their lives. Call this a "contributivist" or "productivist" ethos. Obsolescence is a major threat to the contributivist or productivist ethos because it directly impacts on people's capacities to make those contributions. Furthermore, it's not just the obsolescence itself that creates this crisis. The *impending threat* of obsolescence is also a problem because it makes people aware that the opportunity to make a valuable contribution is dwindling, and so encourages them to compete with one another for those few remaining opportunities. This dynamic of dwindling opportunities and increased competitiveness lies at the core of many of Sparrow's fears of obsolescence.

But, to state the obvious, not all external demands and standards are true sources of value and meaning, and hence not all competitions to meet those demands are worthy of respect. Sometimes our societies demand things of us that are counter-productive to living a valuable and meaningful life. Consider an extreme example. Imagine a physically fit young man living in Nazi Germany. This young man would probably feel a crushing weight of expectation building up around him with respect to making a contribution to that society's militaristic and genocidal aims. Even if he himself did not agree with those aims, or had a strong desire to do something else, he may succumb to those aims due to the

mounting social pressure. Contrariwise, a less physically fit person, or an older person past the prime of youth, may have the freedom to resist the militaristic pressures. Why? Because they would be seen as obsolete relative to the militaristic and genocidal aims. In other words, what we have here is a case in which being obsolete could actually be a net advantage with respect to the desire to live a valuable and meaningful life (this assumes, for sake of argument, these obsolete people were not perceived as enemies to the fascistic project and targeted for extermination).

The Nazi example is instructive. It shows us that not all contributivist demands are sources of value and hence not all forms of obsolescence with respect to those demands are undesirable. But the Nazi example is appealing only because of its historical uniqueness. It tells us that if you live in a society that is in the grip of a dangerous or malevolent ideology, then being perceived (or actually being) obsolete with respect to the ends that this ideology serves might be a good thing (if -- and it could be a big if -- that the society treats its obsolete members humanely and leaves them to their own devices). Surely that is the exception rather than the rule? Surely the norm is that societies are not in the grip of malevolent ideologies and so being able to make a contribution is net positive (and hence being obsolete is not)?

There are two things to be said in response to this. First, many societies are in the grip of ideologies that impose unwelcome, if not outrightly malevolent, standards on most people. For example, standards of beauty, intelligence, physical fitness and so. Consider, once again, the contributivist or productivist ethos that prevails in the modern economic world. This ethos puts a lot of pressure on people to make themselves employable; to acquire levels of education and skills they may not desire; to pursue careers they may not want; to show extreme and sometimes debilitating levels of dedication to their work. This might be acceptable if people could simply reject these external standards and opt out of economic life, but this is not a viable option in the modern world. For the majority of people abiding by the contributivist ethos is effectively compulsory. They have to work in order to survive. If they do not work, when they are able to do so, they are penalised. David Frayne, in his excellent

ethnography *The Refusal of Work* (2015), documents this phenomenon, highlighting both the debilitating effects of a work-obsessed culture on individuals, as well as the difficulties they face if they try to opt out. Others have offered a litany of arguments against the ideology of work (Danaher 2019a; Srnicek and Williams 2015; Bastani 2018). Being relieved from the contributivist standards of the modern labour market, which is pervasive in modern societies, could consequently be a great relief to many people.

Related to this there is the fact that some prevalent external standards are particularly problematic for specific groups of people. Consider for example, the erotic expectations and standards that are imposed on women in patriarchal societies. Under the norms of those societies, they are expected to behave in certain ways and achieve certain standards of beauty and sexual desirability. Writing in the 18th Century, Marie Thérèse de Lambert highlighted this problem and noted that becoming sexually obsolete – as a result of old age – could be liberating for women because it could free them from the yoke of these patriarchal standards. As she put it:

> *[Old age] liberates us from the tyranny of opinion. When one is young, one only dreams of living in the idea of someone else; one must establish one's reputation and create for oneself an honorable place in the imagination of others, and be happy even in their idea; our happiness is not at all real, it is not ourselves whom we consult but others.*

(Quoted in Blank 2018, 299-300)

Lambert is, admittedly, speaking to a bygone age. The patriarchal standards imposed on women are no doubt less severe than they once were (though they are certainly still present). Nevertheless, her observation resonates, in part because it is generalisable. She is highlighting the fact that being freed from the tyranny of other people's opinions and standards can be a good thing if it allows us to pursue a more authentic form of happiness.

This brings me to the second point I wished to make in response to concerns about the Nazi example. Even if we live in societies in which the

external standards are generally beneficent, the mere fact that they are imposed on us at all is contrary to freedom and autonomy. We don't get to choose our own ideals. Actual or perceived general obsolescence could be a great boon for freedom and autonomy in this regard. Because obsolescence implies that external standards are no longer thought to apply to your case, you can decide what to do for yourself. Brian O'Connor (2018) captures the advantage of this lifestyle rather well in his book on *Idleness* when he notes that one of the problems with contemporary society is that we cannot be truly autonomous because we always have to contend with external demands. One way that this manifests itself is in the rejection of idleness as an appropriate state for an adult human being to be in. To be idle is to be frivolous, at best, and a social parasite at worst. Contrary to this, O'Connor suggests that embracing idleness may be the way to secure an idealised form of individual autonomy. This is difficult in a world with rampant contributivist pressures. But this, then, is a reason to welcome widespread, technologically-induced obsolescence. It will enable us to live in a guilt-free state of idleness.

There is an obvious objection to this line of reasoning: not all standards are imposed from the outside-in. Surely, it could be argued, some people develop their own internal standards and then are tragically frustrated and thwarted by their technologically-induced obsolescence when they fail to match them in the desired way. Suppose, for example, Albert has wanted to be a surgeon from his early teens. He sees this goal as the ultimate expression of his own desires and the ideal way to flourish given his competencies and capacities. He spends twenty years of his life educating himself and honing his surgical skills. Sadly, just at the end of his education, a surgical robot is created that can vastly outperform Albert on all relevant metrics. Wouldn't this be a case in which a performance standard was chosen by the individual and in which it was bad thing that they were rendered obsolete by technology?

Albert's fate is certainly a bad thing (for him), but we should question the assumptions underlying the thought experiment. For starters, it is unlikely that Albert's chosen performance standard (skill in surgery) was completely of his own making. He probably picked from among the career options made available

to him in his given socio-cultural milieu. His choice among those options may have been his own – we can ignore the complexities of free will and determinism here – but it was still a choice from among a selection of socially acceptable and educationally feasible options. Furthermore, Albert's tragedy is probably compounded by the fact that he wasn't given other options to choose from. He no doubt grew up – as we all do – in cultures in which we imbibe the belief that success in a career is the primary route to happiness and fulfillment. What if he were given other options? What if he grew up in a culture in which the contributivist ethos was not dominant?

This is where the second argument in favour of obsolescence comes into play. This argument claims that building *a culture of obsolescence* – that is: a culture in which humans are not expected or encouraged or demanded to meet certain external standards, nor valued primarily for their ability to contribute economically or otherwise to society, but are instead freed from the yoke of external standards – might be the best way to cope with rapid technologically-induced change. There are two key assumptions at play in this argument. The first is that advances in technology – particularly in technologies of automation and AI – can be beneficial to human society. They can enable faster, more efficient and more widely distributed goods and services. This isn't always true, of course. Sometimes the technologies can misfire, or be less efficient or less fair than we initially presumed, but if they live up to their potential they can be beneficial. The second assumption is that the major downside to these technologies 'living up to' their potential is that they replace and render humans obsolete in the process. This, however, is only a major problem if we live in a society in which human obsolescence is viewed as a tragedy. Thus, if we build a culture in which obsolescence is welcomed, we can avoid the downside risks of technological change, while embracing its benefits. This is true irrespective of whether technologically-induced obsolescence is inevitable, but the argument might be even more persuasive if obsolesence is inevitable. In that case, we would have the choice of either embracing our inevitable obsolescence or raging against it in futility. The former would seem obviously more desirable than the latter.

There are some obvious objections to this second argument. One objection is that creating a culture of obsolescence is impossible. Maybe we are too wedded to the contributivist ethos to allow for such a systematic and radical change in our governing ideology to take hold? There is undoubtedly some truth to this objection but there is also some reason for optimism. As mentioned earlier in this article, there is an almost endless set of human activities and some human activities are not valued primarily or solely for their instrumental purposes or ends. Think, again, of games or relationships or the bodily pleasures of food, drink and sex. A culture of obsolescence is one in which these kinds of activities, as opposed to contributivist ones, are given pride of place. Furthermore, most activities have several different 'loci of value'. In other words, they are valued for both internal or intrinsic reasons – e.g. because of the skilled or pleasurable performances they involve – and for instrumental reasons – e.g. because they make the world a better place or contribute to economic ends. Building a culture of obsolescence requires shifting focus to away from the instrumental reasons for valuing activities and focusing more attention on the intrinsic reasons for valuing them. This will be a challenge but it is not infeasible and, to a large extent, we can continue to do the things we already find pleasurable; it's just that we derive value from them in a different way.

Another obvious objection is that there is something hollow or sad about a culture of obsolescence. At the heart of any such culture would be the acceptance of the fact that humans can no longer compete or achieve certain standards. This would imply stepping back from the world and retreating into ourselves and our own preferred conception of the good life. This might be a boon for radical liberals or libertarians, but many people will see it as an empty ideal. They want to make a contribution to their societies and to humanity. This is an important objection but one with less bite than might be supposed. For one thing, aspects of the good life that are currently highly valued by many people are still possible in a world in which activities are valued for themselves and not for their broader consequences. Dedication to mastering an non-instrumental, relatively activity (e.g. a game like activity) can allow people to achieve mastery over a certain set of skills and hone their agency. If the activity is one that

requires communication and coordination with others, this can allow for the goods of cooperation and social bonding to flourish.

This connects to a deeper philosophical point that can be made about the nature of human goods and the development of virtuous character. As Alisdair MacIntyre (2007) observes in his influential work on virtue, human goods and virtues are not best defined or understood in the abstract. They only really take on meaning in the context of particular practices or activities.[7] Every activity has goods that are internal to the performance of that activity and this, in turn, leads to a conception of virtuous activity that is tied to those internal goods. Thus, for example, there are goods associated with the practice of chess-playing – strategic playing, intelligent move-making, efficiency – and virtues that a good chess player will develop – decisiveness, courageousness, generosity and graciousness – that can only be defined and understood relative to the rules of the game. This is true for all other goods and virtues too, even if they are not explicitly defined by a set of rules. In other words, there is a way to be a good ballet dancer or baker or bank manager, but there is no way to be good *simpliciter* or *in general*. What it means to be good must be determined and evaluated relative to the nature of particular activities. This theory implies a robust form of pluralism when it comes to figuring out what is valuable and what it means to live a good life. You can pick any activity and figure out a way to perform that activity that achieves certain goods (internal to that activity) and this allows you to build certain virtues (relative to those internal goods). This theory of the good life provides succour to anyone, like me, who wants to argue that a robust form of the good life is possible in a world of technologically-induced obsolescence. It suggests that, far from being hollow or empty, an obsolescent life can be rich and valuable as long as we pursue activities in the right way.

[7] MacIntyre's theory is subtle and complex. He maintains that there are general abstract virtues that can be shared across activities, but he argues that they are relatively hollow and uninformative when considered in the abstract. For an excellent exposition of his theory, which I relied upon heavily for this discussion, see Higgins (2011). Shannon Vallor, in *Technology and the Virtues* applies MacIntyre's theory specifically to developing a virtue ethics for technology. In a sense, the argument in this paper could be allied to Vallor's project insofar as I am suggesting that certain forms of technology could warrant a shift to an ethic of obsolescence. Vallor may or may not welcome this association.

Of course, it is not that simple. MacIntyre's theory is just one among many. Furthermore, MacIntyre himself acknowledges that you can only push this value pluralism so far. Although goods and virtues take on meaning only relative to activities, those activities have to both (a) fit or make sense within an individual's life narrative and (b) cohere with the broader moral tradition/community in which the individual lives. If I steadfastly believe that my life only has value if I make some economic contribution to my society, and if the society in which I live reinforces that belief, then I am unlikely to achieve the good life in a set of activities that confirm my obsolescence. But this is where the idea of building a culture of obsolescence comes into play. If we build such a culture then an individual can structure a life narrative that accepts obsolescence and can have that life narrative reinforced by the community in which they live. This can allow the internal goods and associated virtues of obsolescent activity to thrive.

6. Conclusion

This article has presented an extended analysis of the axiological impact of human obsolescence. It has argued that human obsolescence arises when there is some mismatch between an individual's capacities and some external performance benchmark or standard such that the individual is unable to attain that standard. It has argued that human obsolescence can take at least four distinct forms: actual-general, actual-narrow, perceived-general and perceived-narrow. Of these four forms of obsolescence, the last three are plausible threats to the value of a human life. The first – actual-general – is not plausible except in certain special circumstances (e.g. catastophic individual loss of capacity through illness or injury). Turning its attention to technologically-induced obsolescence, the article has also argued that existing critiques of this phenomenon, particularly Robert Sparrow's critiques of enhancement-induced obsolescence, are not as strong as they first appear, and that there are, in fact, two reasons to embrace general/broad forms of technologically-induced obsolescence: (i) being released from the burden of external standards can be liberating and

empowering, particularly when those standards are bad or unjust and (ii) building a culture that celebrates human obsolescence may be the best way to benefit from technological progress.

References

Agar, Nicholas (2013) *Truly Human Enhancement*. Cambridge, MA: MIT Press.

Agar, Nicholas (2019) *How to be Human in the Digital Economy*. Cambridge, MA: MIT Press.

Baars, Jan (2012). *Aging and the Art of Living*. Baltimore, MD: Johns Hopkins' Universty Press.

Baars, Jan (2017). Aging: Learning to Live a Finite Life. *The Gerontologist*, 57(5): 969–976

Bastani, Aaron (2018). *Fully Automated Luxury Communism*. London: Verso Books.

Blank, Andreas (2018). Anne-Thérèse de Lambert on Aging and Self-Esteem. *Hypatia* 33(2): 289-304

Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age*. New York: WW Norton and Co.

Currier, Richard (2015). *Unbound: How Eight Technologies Made us Human and Brought Our World to the Brink*. New York: Arcade Publishing.

Danaher, John (2019a). *Automation and Utopia*. Cambridge, MA: Harvard University Press.

Danaher, John (2019b). The Rise of the Robots and the Crisis of Moral Patiency. *AI and Society* 34(1): 129-136

De Beauvoir, Simone (1996). *The Coming of Age*. Translated by Patrick O'Briain. New York: WW Norton and Co – originally published in French 1970.

Erdi, Péter (2019). *Ranking: The Unwritten Rules of the Social Game we All Play*. Oxford: OUP.

Frayne, David (2015). *The Refusal of Work*. London: ZED Books.

Frey, Carl Benedickt (2019). *The Technology Trap*. Princeton, NJ: Princeton University Press.

Ford, Martin (2015). *The Rise of the Robots*. New York: Basic Books.

Habermas, Jurgen (2003). *The Future of Human Nature.* Polity Press.

Harriott, Howard (2006). Old Age, Successful Ageing and the Problem of Significance. *Ethical Perspectives* 13(1): 117-141

Higgins, Chris. (2011). *The Good Life of Teaching: an Ethics of Professional Practice*. Oxford: Wiley-Blackwell.

Macintyre, Alisdair (2007). *After Virtue: a Study in Moral Theory (3rd Edn)*. South Bend, IN: University of Notre Dame Press.

Nelson, Mallory Kay, Ashley Shew, and Bethany Stevens. 2019. "Transmobility: Possibilities in Cyborg (Cripborg) Bodies." Catalyst, Special Issue on Crip Technoscience. 5(1)

Mason, Andrew (1990). Nozick on Self-Esteem. *Journal of Applied Philosophy* 7(1): 91-98.

Nagel, Saskia (2010). Too much of a good thing? Enhancement and the burden of self-determination. *Neuroethics* 3(2): 109-119.

Nozick, Robert (1974). *Anarchy, State and Utopia*. New York: Basic Books.

Nyholm, Sven and Smids, Jilles (2018). Automated cars meet human drivers: responsible human-robot coordination and the ethics of mixed traffic. *Ethics and Information Technology* DOI: 10.1007/s10676-018-9445-9

Nyholm, Sven (2020). *Humans and Robots: Ethics, Agency and Anthropomorphism*. London: Roman and Littlefield.

O'Connor, Brian (2018). *Idleness*. Princeton, NJ: Princeton University Press.

Sandel, Michael (2007). *The Case Against Perfection*. Cambridge, MA: Harvard University Press.

Shew, Ashley 2018. "Different Ways of Moving Through the World." *Logic Magazine*, Issue 5: Failure. URL: https://logicmag.io/05-failure/;

Shew, Ashley, (2019) "Stop Depicting Technology As Redeeming Disabled People" *Nursing Clio* 23 April 2019, available at https://nursingclio.org/2019/04/23/stop-depicting-technology-as-redeeming-disabled-people/

Sparrow, Robert (2015). Enhancement and Obsolescence: Avoiding an "Enhanced Rat Race". *Kennedy Institute of Ethics Journal* 25(3): 231-260

Sparrow, Robert (2019). Yesterday's Child: How Gene Editing for Enhancement Will Produce Obsolescence—and Why It Matters. *American Journal of Bioethics* 19(7): 6-15.

Srnicek, Nick and Williams, Alexander (2015). *Inventing the Future*. London: Verso.

Toffler, Alvin and Toffler, Heidi (1970). *Future Shock*. New York: Random House.

Vallor, Shannon (2016). *Technology and the Virtues*. Oxford: OUP.