Life Has Value, but Non-Existence Does Not: In Defense of the Epicurean Reconciliation Strategy

a thesis submitted in partial fulfillment of the requirements
for award of the degree of BPhil in philosophy

Rhys Southan

Abstract

Epicurus argued that death is not bad for the one who dies, because death is annihilation of the experiencing subject and we need a subject of badness for anything to be bad. Epicurus' argument is widely thought to lead to indifference about death, as well as indifference about creating new life, and thus is mostly rejected.

In his essay “Death and the Value of Life,” McMahan offered an alternative framing of the Epicurean position on death. He called this the “Epicurean Reconciliation Strategy” (ERS). The ERS entails that non-existence is not intrinsically good or bad, and non-existence is not comparatively better or worse than existence, but that existence can be non-comparatively bad or non-comparatively good. This non-comparative value provides axiological reasons to create and extend good lives, and to avoid creating and avoid extending bad lives.

In this thesis, I will defend the premises that underlie the ERS, before defending the ERS more directly. I will look at implications that the ERS has on the value of creating and ending lives, and defend it from objections.

0. Introduction 1

1. **Chapter 1: The Symmetric Incomparability of Existence and Non-Existence** 4

    1.1. The Impersonal Valuelessness of Non-Existence 5
    1.2. The Personal Incomparability of Existence and Non-Existence 9
    1.3. The Incomparability of Shorter and Longer Lives 13

2. **Chapter 2: Non-Comparative Value and the ERS** 25

    2.1. Good and Bad Do Not Reduce to Their Comparatives Better and Worse 26
    2.2. Non-Comparative Value is Action Guiding 37
    2.3. Choosing Between Comparative and Non-Comparative Value 39
    2.4. The Epicurean Reconciliation Strategy 49

3. **Chapter 3: Defending the ERS from the Replaceability Objection** 57

    3.1. What Replaceability Means, and Why ERS Entails it 57
    3.2. Cohen's and Dworkin's Axiological Objections to Replaceability 65
    3.3. Accumulating Worseness and Axiologically Irreplaceable Mediocrity 69
    3.4. Velleman et. al.'s Non-Axiological Objection to Replaceability 78
    3.5. Avoiding Replaceability Without Abandoning ERS 81

4. Conclusion 93

5. Works Cited 97

Introduction

Can it be comparatively better or worse for someone that they come into existence? In this thesis, I will join Bader (Unpublished Manuscript A), Bykvist (2007), and McMahan (1981) in answering "No" — defying affirmative answers from Roberts (2003), Holtug (2001), and Arrhenius & Rabinowicz (2015).

Can it be comparatively better or worse for someone that they continue existing as opposed to ceasing to exist? Can it be good, bad, better or worse for someone to cease to exist rather than continue existing? I will join Epicurus (1926 [around 290 BCE]), Lucretius (1910 [around 50 BCE]), Taylor (2012), Suits (2001), and Rosenbaum (1986) in answering "No" to both of these questions — contradicting the affirmative answers from Bader (Unpublished Manuscript), Feit (2002), Nagel (1970), McMahan (1988), Bradley (2009), Luper (2009), and many others.

Can it be non-comparatively good or bad for someone that they come into existence, as well as that they continue to exist? I will join McMahan (1981:105), Parfit (1984) and Bykvist (2007) in answering "Yes" — rebuking the negative answer from Broome (1999) and Bader (Unpublished Manuscript A).

While most of the questions I will be considering here have been thoroughly discussed since Nagel reinvigorated the philosophical debate over death with "Death" (1970), the particular combination of answers I give to these questions leads to a position that has not received much attention. In (1988), McMahan proposed but did not endorse this position, which he called the Epicurean Reconciliation Strategy (ERS). According to ERS, it is not better or worse to exist rather than to never exist. It is also

not good, bad, better or worse to cease to exist. However, there is non-comparative good and bad in coming into existence and in continuing to exist, and this non-comparative value provides axiological reasons to create or not create lives, to extend lives, and to shorten lives.

My original contribution to the debate over the value of coming into existence, existing, and ceasing to exist will be to develop and argue for what I believe are the premises underlying ERS, and to defend ERS as a coherent and compelling wide person-affecting view with relevance to population ethics.

Chapter 1 is devoted to showing that there is no comparative value in coming into existence or in continuing to exist, and there is no value whatsoever in never existing and in ceasing to exist. That is, it is neither better nor worse to come into existence, and it is neither better, worse, good nor bad to cease to exist.

In Chapter 2, I will argue that the lack of comparative value in coming into existence and in continuing to exist does not mean that there is no value in creating new lives or in extending existing lives. This is because it can be non-comparatively good or bad to come into existence and stay in existence. The combination of views I argue for in Chapters 1 and 2 lead to ERS.

The most troubling objection to ERS may be that it entails the "replaceability" of all beings, including rational ones. This means, for instance, that there is no loss in value if everyone in the world were to painlessly die and instantly be replaced by their equally happy, newly created, twins. Chapter 3 will be devoted to refuting the charge that we ought to abandon ERS because of this implication. As a purely axiological

theory, ERS does indeed entail replaceability. But I will show that anti-replaceability is an implausible position to take within axiology. I will also propose a deontological complement to ERS that allows us to avoid replaceability in moral practice. If my arguments are successful, we will have a wide person-affecting axiology that allows us to see value in the creation and extension of lives, and no value in death, and which will not morally permit us to destroy existing lives in order to create slightly better lives in their place.

Chapter 1. The Symmetric Incomparability of Existence and Nonexistence

In this thesis, I will assume a wide person-affecting view in which all moral value is value for sentient beings. A byproduct of this assumption is that I reject the impersonalist view in population ethics which seeks to maximize "impersonal moral value" that is "better simpliciter" or "better for the universe" (Kraut:2012) without necessarily being better or worse for sentient creatures.

Unlike impersonalism, person-affecting views need to address whether or how nonexistence and existence are comparable. In this chapter, I will argue that nonexistence and existence are incomparable in personal terms, and that this incomparability is symmetric. That is, in terms of comparative value, we cannot compare coming into existence as opposed to never existing – a relatively uncontroversial but disputed claim[1] – and we also cannot compare continuing to exist as opposed to ceasing to exist (a highly controversial and widely disputed claim[2]).

In section 1.1, I will describe my conceptualization of nonexistence, with the aim of establishing the valuelessness of nonexistence, and of showing why the notion of "possible beings" can be misleading – even if I am not able to fully avoid such terminology myself.

[1] Bykvist dubs many of those who defend the personal comparability of existence and nonexistence the "Scandinavian existentialists." These include Arrhenius & Rabinowicz (2015), Johansson (2010), and Holtug (2001). Roberts (2003) is a non-Scandinavian sympathizer.

[2] Epicurus is the most famous (apparent) proponent of this view. In the current philosophical debate about the value of death, those who disagree with Epicurus seem to outnumber those who agree.

Section 1.2 will cover more familiar territory. I will summarize what Bader, Broome and McMahan have said against the personal value comparability of nonexistence and existence. I will also offer some thoughts of my own.

In 1.3, I will argue that the personal incomparability of existence and nonexistence is symmetric: that is, it applies to comparisons between never existing and coming into existence, and also to comparisons between continuing to exist and ceasing to exist. If this is right, then those who say we should be axiologically neutral about creating lives also must accept axiological neutrality about how long existing lives last. If my arguments are sound, they provide a powerful reason to reject views claiming that there is no value in creating new lives. In addition, my arguments for the symmetric valuelessness of nonexistence, and the symmetric incomparability of existence and nonexistence, motivate the view I will defend in the second chapter – ERS.

1.1 The Impersonal Valuelessness of Nonexistence

Though they have much to say about the value of existence contrasted with nonexistence, philosophers of death and population ethicists typically have not said much about nonexistence itself, beyond referring to it as "a blank" (Nagel (1970:73), Timmerman (2015:31), Broome (2004:13)). The quietism about nonexistence may be appropriate. It is, after all, nothing. However, in this section, I will attempt to say more. My intention in conceptualizing nonexistence is to more solidly ground my claim that nonexistence truly is valueless, and symmetrically so – whether it is the nonexistence of never existing or the nonexistence of having ceased to exist.

I define sentient existence as the capacity for experience, consciousness, or awareness. If there is no afterlife, as I will assume there is not, death is the annihilation of a being's sentience and the negation of one's existence. When a sentient being ceases to exist, that capacity for experience, consciousness or awareness disappears.

We often refer to "nonexistence," with reference to possible beings who never exist, or once-existing beings who cease to exist. This may be unavoidable, and at times helpful, but it can also be misleading. There are neither non-existent beings before lives begin nor non-existent beings once lives end. We cannot accurately speak of anyone's nonexistence, because nonexistence refers to either the destruction of a being or the never coming to be of a being, and these are necessarily impersonal. (That is, there is no personal point of view in nonexistence.) Creating someone does not bring this someone out of nonexistence, and destroying someone does not return this someone to nonexistence. Before someone's existence, there was no someone to be brought into our world, and after someone is destroyed, there is no someone to be returned anywhere.

We can create more sentient existence, and we can do the opposite by failing to create sentient beings or by destroying sentient beings so that there is less sentient existence. But we cannot create more nonexistence, and we cannot bring about less nonexistence. Nonexistence does not grow as existence contracts. Sentient existence grows and contracts without any effect on its alleged opponent.

There is existence of varying lengths, but no length of nonexistence. One does not have a longer or a shorter nonexistence depending on when one ceases to exist, or whether one never exists, and one does not have an infinite nonexistence either way.

Nonexistence is not had by anyone. Only existence can be had. If nonexistence is anything at all, it is the borders of existence.

None of us has experienced nonexistence, because nonexistence cannot be experienced – it entails the absence of experience. It would not make sense to say, "It is a very different experience to have experiences than to not have experiences," since not having experiences is not an experience. It cannot be different for me that there is a me instead of there not being a me, since there is nothing for me to contrast my own existence with. None of us has ever been non-existent, or ever will be non-existent, as nonexistence is mutually exclusive with being. There is nothing it is like to not exist.

As I said, philosophers sometimes refer to nonexistence as "a blank." There is no value in a blank. The blank of never existing is valueless, and so is the blank of having ceased to exist.

This last claim perhaps contradicts the common view that death can be either good or bad; as the negation of life, death is commonly thought to be good when life is bad and thought to be bad when life is good. This implies value. However, we should not assume that the negation of something is necessarily its opposite.

Kazez makes a point like this with a hypothetical example in which her family goes to Paris (2011). While in Paris, they have good experiences of the taste of croissants, and the pleasure of seeing real Renoirs. There are also bad experiences such as pain when gravel gets into their shoes. And then there are pains that are absent from Paris, such as the smell of an Arkansas pig farm or the London tube, and the lack of an Indonesian monsoon. Do these latter absences that apparently have nothing to do with

Paris get counted amongst the goods of Paris as well? Kazez says no: “The *bads* in London, or Arkansas, or Indonesia, are points in favor of going to Paris, but they aren't *goods* in Paris.” An anonymous commenter on her post adds, “Suppose I refrain from writing a piece of poetry and I am a very bad poet. Does not having written my ugly poem make for something beautiful?”

If absences could be good or bad, there would be a diverse range of rankable goods and bads in a barren universe. The absence of eternal suffering in an empty universe would be very good, whereas the absence of itches would be good to a lesser extent. The absence of eternal pleasure would be very bad, whereas the absence of lollypop licks would be bad to a lesser extent. Depending on how these goods and bads of absence weighed against each other, a barren universe could be either good or bad overall, despite having absolutely nothing in it and despite its apparent neutrality. I think this shows we should doubt that there are goods or bads in absences. Value resides in things; it cannot reside in absent things, because there are no absent things.

Another reason to doubt the value of absence is that it seems different to replace a pain with pleasure than to merely remove a pain. Removing a pain is by itself a negation or neutralizing of that pain – though removing a pain often does provide positive pleasure in the form of relief. If the mere negation of pain itself is good, it is not good in the same way that replacing a pain with pleasure is, and the difference between these two is not that the pleasure is “more good” than the reduction of pain is. Pleasure is the presence of something positive, whereas the removal of pain leads to the absence of something negative.

I therefore claim that nonexistence is best conceptualized as having no value whatsoever.

This means that we cannot call someone's existence better or worse for them. Better and worse are comparative terms referring to counterfactual improvements or reductions in value. However, before there can be an improvement or reduction in value for someone, there needs to be that someone. Value that comes into existence because the bearer of that value comes into existence cannot be an improvement or reduction for the value bearer, since there is no pre-existence value level to measure the later value against.

In the next section, I will look at two general attempts by philosophers to establish – in contradiction to my conclusion here – that it can be better or worse to come into existence. I will summarize and endorse counterarguments by Bykvist, Bader, McMahan and Broome that suggest these attempts fail.

1.2 The Personal Incomparability of Existence and Nonexistence

Not everyone agrees that nonexistence is valueless. Some philosophers claim that nonexistence has a value of zero wellbeing for a potential individual, which can be coherently contrasted with the positive and negative wellbeing scores we find in existence (Roberts:2003). If a life's wellbeing would be overall positive, then it would be better to create the life, because creation would bring the wellbeing score from zero to something higher (Holtug:2001). If a life's wellbeing would be overall negative, then it would be worse to create the life, because that would dip the wellbeing below neutral.

If this is right, then there certainly is value in creating new lives.

But this approach is widely deemed unacceptable. Reporting on personal correspondence with Broome, Bykvist writes, “Broome has objected that having no value is not the same thing as having zero value. For instance, logic has no temperature, but that does not mean that it has zero temperature and thus is colder than the ocean” (2007:343). If this is a good analogy, it is more plausible to say that nonexistence has no wellbeing instead of zero wellbeing. Yet we cannot compare some level of wellbeing with no wellbeing whatsoever.

I see two ways for the advocate of comparing existence and nonexistence (“the comparativist”) to attack this temperature analogy. One is to note that we can create wellbeing from materials that have no wellbeing by transforming these into a sentient creature, whereas creating temperature from a non-physical concept like logic is inconceivable. This suggests a disanalogy. However, it does not seem to be an important one, and so I do not think this claim helps the comparativist. The comparativist might then want to say that while logic is not colder than the ocean, we can at least say the ocean has “more temperature” than logic does, since the ocean has temperature and logic does not, and some temperature is more temperature than no temperature. Bringing this back to wellbeing, we could say that the creation of a sentient creature leads to there being “more wellbeing” than the non-creation of a sentient creature, which suggests some numerical comparison of wellbeing is possible here.

This latter point seems correct, but Bykvist convincingly claims it is a trivial truth. Responding to Roberts, he argues that we should not confuse the existence of “more

wellbeing" due to the creation of someone with this new individual being better off for having come into existence (Unpublished Manuscript:13-14). The comparability problem is not that we cannot say that the existence of wellbeing is in some way different than there existing no wellbeing. As Broome, McMahan and Bader have said in different ways, the problem is that we cannot establish a "for whom" existence is better for when the alternative is a subject's absence. Broome (1999:168) and McMahan (2008) put it this way: for something to be better (worse) for someone, the alternative needs to have been worse (better) for them. But there is no one for whom never existing is worse (better), and therefore coming to exist cannot be better (worse).

Bader has a subtly different explanation which is more in line with what Bykvist (2015) has argued. Bader writes, "The first problem for comparativism is that in the case of nonexistence there is a missing relatum of the betterness relation, making it impossible for the relation to connect a case of existence with one of nonexistence" (Unpublished Manuscript A:11). The betterness relation in question is the "betterness of lives" relation. If P is better off in A (with A being a world in which P exists) than in B (with B being a world in which P does not exist), then this means that the life-betterness relation holds between two lives L1 and L2, each of which are lives of P, which are in A and B respectively. But then this suggests there is a life of P in B (L2), so P exists in B, which contradicts our assumption that P does not exist in B.

Attempting to salvage what might be called "full comparativism" does not look promising, but there is a milder variant of comparitivism to consider. Bader it calls "restricted comparativism." In this view, it is not worse (better) for the never existent

that they never exist, but once someone exists, we can say that it is better (worse) for them that they came into existence, and that it would be worse (better) for them to cease to exist (Arrhenius & Rabinowicz:2015). However, once they cease to exist, it is not worse (better) for them that they have ceased to exist, since there is now no one for whom it is worse (better).[3]

Both Bykvist (Unpublished Manuscript) and Bader (Unpublished Manuscript A) have criticized this view. Bader says it runs afoul of a commitment to axiological invariance. He writes, "the betterness ordering cannot vary depending on the world from which things are evaluated" (Unpublished Manuscript A:15). Bykvist summarizes his own criticism, "if a state of affairs cannot contribute any value for you, then it cannot have any value for you. In its most general formulation the idea is that if a property or state of affairs cannot contribute any value to the things that exemplify it, then it cannot have any value" (2015:92). Bykvist also offers a deductive argument against restricted comparativism that I will not summarize here (Unpublished Manuscript). While there may be more to say in defense of restricted comparativism, the fact that it arrives at different judgements about the value of nonexistence depending on whether one exists or not is inventive but also metaphysically suspicious**,** and it clashes with what I see as the more plausible claim, that nonexistence is always valueless. I believe the criticisms from Bader and Bykvist are sufficiently convincing that we should treat existence and nonexistence as fully incomparable in terms of

[3] This last claim about ceasing to exist may not be a necessary component of this view, but this is the version of the view that I am interested in here.

personal value. It cannot be better or worse to come into existence than to never exist.

1.3 The Incomparability of Shorter and Longer Lives

If coming into existence is not comparatively valuable, then perhaps there is no value whatsoever in coming into existence. This is what Bader argues (Unpublished Manuscript A). He arrives at this conclusion by combining three claims:

1. All moral value is reduceable to comparative value (better and worse) (Unpublished Manuscript B:6-7).
2. All moral value is personal value – it is value for beings.
3. Existence and nonexistence are incomparable in personal terms.

If existence and nonexistence are incomparable in personal terms (3), then it is not better or worse for someone that they come into existence. If anything that can be called value can also be called better or worse (1), then there is no personal value in coming into existence. If all value is reduceable to personal value (2), then there is no value in coming into existence.

I accept 2 and 3. In the next chapter, I will argue that we should reject 1. For the rest of this chapter, I will defend a claim that would make the combination of premises 1-3 too radical for most people to accept. This claim will also help to establish an essential component of the Epicurean Reconciliation Strategy – that there is no value in going out of existence. I will call the claim premise 4:

4. Making a comparative judgment between continuing to exist and ceasing to exist would require comparing existence and nonexistence.

If judging the value of lives of different lengths involves comparisons of existence and nonexistence, then we cannot fully compare lives of different lengths by reference to comparative value, because of (3). This invokes a view like the one Bader calls "temporal part affecting view (TPAV)" whenever we want to compare the value of lives of different lengths (Unpublished Manuscript A:87). Bader writes that if we accept TPAV, "Betterness relations amongst lives would be determined by betterness relations amongst the temporal parts composing those lives. As a result, comparisons of different lives would be subject to an equinumerosity constraint and lives containing different numbers of temporal parts would end up being non-comparable" (88).

If premises 2-4 are right, lives of different lengths are at least partially incomparable in terms of their comparative value, but – I would claim – potentially partially comparable as well. Partial incomparability entails that segments of longer and shorter lives can be compared with each other, even though the lives as wholes cannot be ranked in terms of comparative value. If there are two lives, and both are identical to each other except that one life lasts a second longer, we can make a nearly full comparison of the two lives in terms of comparative value. These lives are identical in value, aside from the extra second in the one life which is incomparable in value with reference to the other life.

Here is an implication of this semi-comparability of longer and shorter lives, given premise 4. We might wonder whether Life A, which is a 100-year life at level-5 goodness, is better than Life B, which is a 50-year life at level-5 goodness. If we accept my premise 4, what we can say is that any particular moment from Life A is comparable

to any particular moment from Life B as equal in value. If we tried to make an overall comparative value judgment, we might say that 50 percent of Life A is comparable with Life B as equal in value, and this can be any 50 percent segment of Life A that we choose. But Life A as a whole cannot be compared to Life B as better or worse because there is not enough Life B for a full comparison.

If we accept 4, "better and worse" refer to the value of the content of existence – in particular, to a gain in value or a loss in value that someone achieves or sustains within their existence – but not to value differences due to variations in length of existence. We can make someone's life better by giving her a gift when she would otherwise have existed without this gift, but we cannot make her life better by making it longer. Similarly, we cannot make her life worse by shortening it. On their own, premises 1-3 are consistent with there being value in the extension of already existing lives. By adding premise 4, we are forced to the conclusion that there is potentially no value in extending an already existing life. That is, we should be axiologically neutral about creating new lives, and we should also be axiologically neutral about extending or shortening existing lives.

If we reject premise 1, we are not forced to accept the valuelessness of extending existing lives, because even if life length cannot affect a life's comparative value – its betterness or worseness – extending existing lives can add non-comparative good and bad. Rejecting premise 1 also entails there being value in the creation of new lives. This strikes me as plausible anyway, and certainly more plausible than denying the value of extending existing lives. However, Bader and others who like premises 1-3 can deny the

valuelessness of extending existing lives by denying 4. This is what Bader does. He claims that once we have a life, the comparative value measure of a life is not just about the value fluctuations that occur within the borders of the life – it is also about how far apart the borders of its existence are. Bader writes:

> [T]here is a point of view, namely the personal point of view, from which a shorter life can be compared with a longer life … This makes a life an axiologically significant unit and allows us to compare a longer life with a shorter life without comparing the temporal parts making up those lives. (Unpublished Manuscript A:87-88)

Even those who do not like premises 1-3 may want to deny my premise 4. McMahan does not accept premises 1 or 2, but he does not agree with 4 either. He writes:

> [Continuing to exist rather than ceasing to exist] may seem to require a comparison between existence and nonexistence. But this is a mistake; one can make the choice between continuing to exist and ceasing to exist by simply comparing two possible lives - a longer one and a shorter one - and asking which would be the better life. This kind of solution is not available when the choice is between life and total nonexistence rather than between continued life and posthumous nonexistence. (1998:215)

Comparative value certainly implies comparison. My disagreement with Bader and McMahan on premise 4 is over what the relevant sort of comparison is when we would compare the value of possible lives of different lengths. Bader and McMahan believe this comparison is between different lives. I will argue instead that the relevant comparison would be between existence and nonexistence.

I am partly inspired by a letter that Epicurus wrote to his friend Menoeceus (1926 [Around 290 BCE]). In his letter, Epicurus essentially asks: How can death be worse for X if there is no X because X has ceased to exist? How can it be better for X to exist longer if the alternative is for X to not exist at all, in which case there may be no second relatum for a comparison? I will try to demonstrate that the answers to both these questions are "it cannot be." This entails that death is neither good, bad, better nor worse.

Bader claimed, "[T]here is a point of view, namely the personal point of view, from which a shorter life can be compared with a longer life." I deny that this is so. A personal point of view entails a "something it is like" – it entails experience. As McMahan, Bader and I all agree, there is no personal point of view when it comes to never existing. Never existing is not an experience. There is nothing it is like to never exist. In contrast, there is something it is like to have a life. We cannot compare a "something it is like" with a "nothing it is like," and so we cannot compare a life with no life. So far Bader, McMahan and I agree. My disagreement with them begins when I add that while there may be something it is like to have a long life, *there is nothing it is like to have a short life*. Certainly, there is some "something it is like"-ness within a

short life. This comes from the life experiences that a particular short life contains. That a short life has experiences and a personal point of view does distinguish a short life from no life. But it does not distinguish a short life from no life in the right sort of way to make a longer life comparable, on the whole, to a shorter life. There is a personal point of view in a short life, but this personal point of view is not *qua* short life; what makes a life short is that the life's personal point of view ends where it does.

I find it plausible that there is something it is like to have a long life. But this does not help us compare a long life and a short life if there is nothing it is like to have a short life. After all, there is something it is like to have a life, but we cannot compare having a life to never existing because there is nothing it is like to never exist. We need a "something it is like" on both sides, or we have a missing, or at least incomplete, relatum. And so what Bader and McMahan need to show is that there is something it is like to have a long life *and* something it is like to have a short life.

There may be something that it is like to have a "short life *so far*." A very young person could know what this is like, or an older person might be able to remember it. Having a short life so far could be the experience of having a relatively new life. Perhaps this is the feeling that the first boundary of a life, the life's beginning, is not too far in the past, and not very much has happened yet. But the feeling of having a short life so far does not tell us what it is like to actually have a short life overall. Someone can have the experience of having a short life so far without having a short life overall. This is because the experience of having a short life so far can be part of a long life. It will just be in the very beginning part of the life.

There is something it is like to sense death around the corner, and this can be combined with the feeling that the beginning border of a life is not too far behind. But this also does not give us what it is like to have a short life. Someone can have the experience of death being around the corner, and of the beginning border of life not being far in the past, and yet be wrong that death is around the corner and go on to live a long life. Even if they are correct that death is around the corner, and that the initial boundary of their life is not far in the past, there is nothing it is like for the second boundary of life to be not merely looming but actually reached. There is nothing it is like for one's life to have been a short life because the temporal distance between the first and second boundary of a life turned out to have been narrow. There is not an experience of a life with extremely contracted boundaries because while these boundaries can be contemplated, they cannot be experienced – they mark where experience does not reach. As I claimed in the first section, there is no experience of having fewer experiences. The "I haven't had a lot of life yet" experience *can* be compared to the "I've had a lot of life," experience, but this cannot be the sort of comparison that Bader and McMahan have in mind. We should not confuse the experience of a life "being short so far" with what it is like for this little bit of life we have had so far to be all the life we get because the life ultimately will be short. There is not something it is like to have had all the life you are going to have, and now it is all over – and this was not very much life.

There is another way to make the point. Imagine there is a button I can press that will end my life. When I contemplate pressing this button, I am not concerned with

whether my pressing this button or not will make the difference between a longer or shorter life for me. I do not decide to press the button because “it is worse for me to have a longer life” or “it is better for me to have a shorter life.” I do not decide against pressing the button because “it is better for me to have a longer life,” or “it is worse for me to have a shorter life.” If I press the button, it is because I would rather not exist anymore. If I do not press the button, it is because I would rather keep existing. If I have a good life and I press the button anyway, doing so does not give me a worse life. It gives me no life. If I have a bad life and I press the button, doing so does not give me a better life. It gives me no life. The question this button poses as I consider pressing it is “to be or not to be?” – not “a longer life or a shorter life?” “To be” is to exist and “not to be” is to not exist; this button is making me choose between existence and nonexistence.

I could have died yesterday. Instead, I continued existing. I might wonder, “How does my having continued to exist compare to the possibility of my having died yesterday?” To me the most natural answer is: “If I had died yesterday, there would be no me having experiences. Instead, because I lived another day, here I am having experiences.” I would not necessarily think, “If I had died yesterday, I would have had a shorter life.”

I could die tomorrow. I might wonder, “How does the possibility of my dying tomorrow compare to the possibility of my continuing to exist tomorrow?” To me the most natural answer seems to be, “If I die tomorrow, there would be no me having experiences. If I do not die tomorrow, there I will be, having experiences.” I would not

necessarily think, “If I die tomorrow, I will have a shorter life,” or “If I do not die tomorrow, I will have a longer life.”

I could have never been born. Instead, I was born. I might wonder, “How does my having been born compare to the possibility of my never having been born?” To me the most natural answer is: “If I had never been born, there would never have been a me having experiences. Instead, because I was born, there has been a me having experiences.” And then I might wonder, “What would it be like if I had never existed?” To me the most obvious answer is, “It would be like nothing.” I might then wonder, “What would it be like if I stopped existing tomorrow?” To me the most obvious answer is, again, “It would be like nothing.” I would not think, “It would be like having a shorter life.”

No matter whether I am thinking about ceasing to exist or never having existed, compared to my existing, it seems the relevant difference is between me existing and having experiences and me not existing and not having experiences. Never existing and ceasing to exist are not relevantly different in this regard.

When philosophers compare longer and shorter lives, they often illustrate the comparison by drawing two lines, one of which is shorter than the other. We look at both of these lines, see that the longer one has more wellbeing due to its extended length, and arrive at the conclusion that it is therefore the better life (assuming it is a good life). But it is not enough for the longer life to have more wellbeing than the shorter life for this kind of comparison to make sense. After all, we could draw a line that is meant to illustrate a life of some quality and accurately say this life has more

wellbeing than no life at all, and erroneously conclude that having this good life is better than never existing. This is what Bykvist criticizes Roberts for doing. If comparing longer and shorter lives is to avoid this problem, we need to show that the overall welfare levels of two lives being different due to length difference is an indication that one of these lives is better off than the other, and one of these lives is worse off than the other. We also need to show that the longer life does not have a higher wellbeing score because it is (in part) being compared with no life at all.

But of course I have been trying to show that the shorter life does not leave someone worse off than the longer life (because a life cannot be experienced as ultimately short – a life cannot be "short for" the person whose life it is or was), and that being partially compared to no life at all is exactly why the longer life has a higher wellbeing score. To re-elaborate on the first point, the never-existent do not experience having fewer experiences. This is because there is never any subject to have fewer experiences. But someone who lives a shorter life also does not experience having fewer experiences than they would have had with a longer life. With both never existing and ceasing to exist, we lack a subject of fewer experiences – we therefore lack a subject of less good; we lack a subject of worseness.

To re-elaborate on the second point, when we decide whether to keep existing or not, we are deciding between existence or no life at all. If we end our life when it could have been longer, and it turns out a longer version of our life would have had more wellbeing, then it seems this possible longer version of our life has this wellbeing advantage because some amount of welfare is more welfare than no welfare at all.

When comparing a longer life to a shorter life, we are in part comparing life to life, and we are in part comparing life to no life.

There is something it is like to have a life of bliss, and there is something it is like to have a life of torture, but there is nothing it is like to have a short life. Someone whose life is worse for having headaches is subjected to those headaches; the worseness of those headaches is part of their life. No one is subjected to a shorter life. The reason a shorter life has a smaller positive welfare score than a longer life is not that there is an increase in bad compared to the other possible life, or a drop in goodness within the boundaries of existence compared to the other possible life, but because welfare has gone from some amount of welfare to no welfare whatsoever, through the absence of the welfare subject. “No welfare” is in some sense less welfare than some welfare, but to paraphrase Bykvist, “no welfare” does not leave one “worse off” – instead, it is the absence of one. A short life is not *short for* the being with the short life, unlike a blissful life that is *blissful for* the being with the life, and unlike a torturous life that is *torturous for* the being with the life. This means a short life is also not *worse for* the being with the short life. And as person-affecters, we should be concerned with the properties of lives that are *for* those lives. A contraction in the borders of a life is not one of these properties.

When a philosopher says that it is worse to die earlier rather than later, all this can mean is that if the life had lasted longer, there would have been good in the future portions of the life, but instead, there is no value at all. This is a consideration of value that could have been. But this does not mean that there is something different about the

lack of wellbeing for a once-existent life and the lack of wellbeing regarding a never-existent life. The problem for those who deny the value of coming into existence is that "value that could be" is value we can foresee producing by creating new lives, in just the same way we can foresee producing it by extending existing lives. We know that if we extend the life of an existing being which would otherwise cease, there will exist someone to get future value when there otherwise would not have been. Yet this is just as true about the creation of beings. If we create a new person, there will exist someone to get future value when there otherwise would not have been. Whatever value we see in extending someone's life, we must also see in creating someone. And whatever value (or lack thereof) we see in never existing, we must also see in ceasing to exist. Since never existing is not good, bad, better or worse, ceasing to exist is not good, bad, better or worse. Value and value bearers do not just rise together. They also fall together.

Chapter 2. Non-Comparative Value and the ERS

I have four main goals with this chapter.

The first is to argue that good and bad do not reduce to their comparatives of better and worse. I will do this in 2.1 by defining good, bad, better and worse in a way that I think best captures what we typically mean by these terms, and which leaves open the possibility of non-comparative value: good and bad. I will then criticize Broome's case (1999) for the reduceability of good and bad to their comparatives of better and worse. If my criticism is apt, I will have made room for the possibility that coming into existence and continuing to exist can be good or bad for one despite it not being better or worse for one.

Second, I will show that this non-comparative value of good and bad is action guiding even when it comes apart from comparative value, as I claim it does with regard to creating new lives or extending or shortening existing lives. In 2.2, I will argue that something's being good for someone without it being better for someone provides a reason to bring it about, while something's being bad for someone without it being worse for someone provides a reason to prevent it from occurring. This will demonstrate that there is value in creating and extending lives, even though this value is not comparative.

One might assume that even if non-comparative value is action guiding, comparative value is more important than non-comparative value whenever the two come apart. In 2.3, I will claim the opposite: that it is non-comparative value rather than

comparative value that we should prioritize. We ought to be at least as concerned with how much goodness or badness for beings there is in an outcome as with what is better or worse for beings in an outcome. If I am persuasive, I will have shown that creating and extending lives can be as valuable as improving lives that will exist no matter what we do.

In 2.4, I will describe where all of my arguments so far have led us: to my version of the view that McMahan called the Epicurean Reconciliation Strategy. According to this view, it is not better or worse to come into existence, but creating and extending lives can be non-comparatively good or bad for those lives. Furthermore, never existing is not good, bad, better or worse, and the same is true for having ceased to exist.

2.1 Good and Bad Do Not Reduce to Their Comparatives Better and Worse

Before I can establish that not all value is comparative, I should clarify what it might mean for value to be either comparative or non-comparative. I will do this by offering some definitions of good, bad, better and worse that are consistent with the person-affecting view that I accept. What I offer are loose, tentative definitions of evaluative terms that provide a sense of what I mean by comparative and non-comparative value.

Good

Desirable, attractive, and/or a positive contribution to a life.

Bad

Undesirable, aversive, and/or a negative contribution to a life.

Better

More good or less bad for someone than an alternative.

Worse

Less good or more bad for someone than an alternative.

Bykvist offers similar definitions to distinguish between comparative and non-comparative value, though instead of referring to non-comparative good or bad, he refers to "absolute benefits" and "absolute harms." He writes, "[A] person enjoys (suffers) an absolute benefit (harm) if and only if she is in state that is good (bad) for her, whereas a person enjoys (suffers) a comparative benefit (harm) if and only if she is in state that is better (worse) for her than some alternative state" (2007:337). We might therefore think of non-comparative value as intrinsic value in contrast to the comparative value of better and worse. Parfit uses intrinsic in this way (2018:131).

My definitions, as well as those of Bykvist, allow comparative value and non-comparative value to occasionally exist independently of each other. For instance, someone's life could have nothing positive (good) or negative (bad) in it, and the life could be better for lacking suffering and worse for lacking pleasure. In the other direction, it could be good or bad for one to come into existence, without our being able to say that coming into existence is better or worse for this individual, since the goods

and bads of her existence are not an increase in value for her compared to the alternative of there being no her at all.

However, good, bad, better and worse can be defined in such a way that comparative value and non-comparative value cannot come apart. Broome tries to accomplish this by defining good as purely relational and therefore always reduceable to better – claiming, in other words, that there is no such thing as non-comparative value or intrinsic value that cannot be described in comparative terms (1999:162-173). To show that there can be non-comparative value in coming into existence even though there is no comparative value in coming into existence, I therefore need to prove that Broome is mistaken.

Here Broome makes his positive argument for the reduceability of good to better:

> Let us say a property F is 'reducible to its comparative' if any statement that refers to the property has the same meaning as another statement that refers to the property's comparative, Fer than, instead. I think goodness is reducible to its comparative. Anything that can be said using 'good' can be said using 'better than'. To put it another way, if you knew everything about betterness – of every pair of alternatives, you knew whether one was better than the other, and which – then you would know everything there is to know about goodness. There is nothing more to goodness than betterness. (1999:164)

Unless Broome has said more about this elsewhere, this appears to be the entirety of his positive argument for the reduceability of good to better. While it feels like this must be begging the question, the claim that coming into existence is good or bad without being better or worse seems to be a direct refutation of Broome's argument[4]. Broome is apparently basing his case on language and meaning. But it is obviously not true that "Anything that can be said using 'good' can be said using 'better than'" if some people want to say it is good to come into existence, but that coming into existence is not better than never coming into existence. Therefore, "good" is not reduceable to "better than."

Another objection to Broome's approach is that it can be linguistically non-intuitive and confusing. Reducing good and bad to their comparatives means that except for the ultimate extremes, everything is both good and bad, including things we might typically want to think of as purely good or purely bad. Even the second worst thing in the world is good because it is better than the very worst thing, and even the second best thing in the world is bad because it is worse than the very best thing. A 10,000-year life of bliss is bad because it is not a 10,001-year life of bliss, and a 10,000-year life of torture is good because it is not a 10,001-year life of torture. At the same time, the second worst thing is also bad because it is worse than almost everything else, and the second best thing is good because it is better than almost everything else. A 10,000-year life of bliss is good because it is better than a 9,999-year life of bliss, and a 10,000-year life of torture is bad because it is worse than a 9,999-year life of torture. But Broome's value theory cannot allow that there may be various degrees of bad things that are less

[4] This point was suggested by [redacted.]

bad than some things but are still just bad, or that there may be various degrees of goods that are less good than some things but are still just good. For Broome, only the very worst thing can be bad without also being good, and only the very best thing can be good without also being bad.

If we contemplate a 10,000-year life of torture when the alternative is a 10,001-year life of torture, we may want to say, "The 10,000-year life of torture is better, but it is still bad." For Broome, this is just the same as saying, "The 10,000-year life of torture is better, but it is still worse" – better than 10,001 years of torture but worse than 9,999 years of torture, for instance. This interpretation does not seem to fully capture what we mean by "better but still bad." A 10,000-year life of torture being bad need not have anything to do with it being worse than a 9,999-year life of torture. It is bad whether we compare it to anything else or not.

Furthermore, we may often want to use "better" when referring to something we think is bad, or "worse" when referring to something we think is good, but unless we are careful in specifying what we mean by this, Broome will be prone to thinking that the "better but bad" thing is "good" and the "worse but good" thing is "bad."

If good and bad are non-comparative, as I claim, then the second best thing in the world is good and it is not bad, but we could call it "less good" than the very best thing. Similarly, the second worst thing in the world is bad and is not good, but we could call it "less bad" than the very worst thing. Though he was not addressing Broome, Smuts suggests language like this for similar reasons (2012).

Broome's approach seems most obviously unsatisfactory when we want to

distinguish between positive states and negative states – a division that comparative-only value does not seem to acknowledge at all. Imagine a chair that can give the person sitting in it various degrees of pains and pleasures. There is a dial for pain, which can go from 0 to 11, and a dial for pleasure, which can go from 0 to 11. Someone sits in the chair, with both dials set to 0. After a few moments of this neutral setting, Broome sets the pain dial to level 6 and leaves the pleasure dial at 0. "That's worse!" the person in the chair screams. "In fact, that's downright bad!" Broome would seemingly be confused by this claim. Given that bad and worse are interchangeable for him, to Broome's ears this should sound like, "That's worse! In fact, that's downright worse!" or "That's bad! In fact, that's downright bad!" For the non-comparative value advocate, in contrast, this statement makes sense. Level-6 pain is worse than level-0 pain, because it is more bad than 0 pain, and level-6 pain is downright bad because it is an aversive sensation.

In short, if we say that good reduces to better, and bad reduces to worse, it becomes difficult or impossible to distinguish between improvements in value and states that are valuable in themselves, and also between diminishments in value and states that are disvaluable in themselves.

Of course, Broome could accuse me of imagining distinctions that do not really exist. He believes that to avoid the reduceability of good to better and bad to worse, we need a zero point for value, with positive and negative degrees: "things that are better than the absolute zero will be absolutely good; things that are worse, absolutely bad" (1999:165). But Broome does not think this zero point for good exists. This is his

negative case against non-comparative value. Broome interprets Parfit (1984) as proposing a zero point for value when Parfit claims that a life which is barely worth living is a life that is slightly better than an empty life. Broome cites Parfit's implied case for this zero point before explaining why he is unconvinced. Broome quotes Parfit, "A life could be [barely worth living] either because it has enough ecstasies to make its agonies seem just worth enduring, or because it is uniformly of poor quality. Let us imagine the lives…to be of this second drabber kind. In each of these lives there is very little happiness" (387). Commenting on Parfit's example, Broome writes, "According to this idea, 'a life of zero goodness' means a life that is neither better nor worse than an empty life. … The natural zeros of the good and bad things in life have given us a naturalistic zero for goodness" (1999:169).

Parfit's allegedly naturalistic zero point for good and bad allows Broome to attack non-comparative value with the classic anti-naturalism argument: G.E. Moore's open question argument (Broome 1999:169-170). So much has been written about Moore's notorious argument that I think it is best if I let Dancy summarize it:

> The best version of Moore's actual argument, in my view, was that there is a difference between two questions:
>
> 1. Is being good being what we desire to desire?
> 2. Is being what we desire to desire being what we desire to desire?

> The first question is 'open' in a way that the second is not. Moore concluded from this that being good is not being what we desire to desire—and then generalized this conclusion to all possible candidate 'definitions' of being good, or of goodness, in natural terms. Sometimes the conclusion is expressed in terms of properties. The property of goodness is not definable in natural terms, because all possible definitions fall to the distinction between open and closed questions; they render closed a question that should be open. (2013:732)

Dancy does not see this as a successful argument, and cites Frege and Kripke as both indirectly supplying refutations. The Frege-inspired refutation is more straightforward. In Fregean terms, the open question argument merely shows that "is good" is a different presentation of a concept than "is what we desire to desire." Moore has not, however, shown that "is good" and "is what we desire to desire" refer to different concepts (732).

I find Dancy's Fregean criticism of the open question argument persuasive. But I think Broome's use of Moore's argument can be attacked even more thoroughly.

Broome says that Parfit's zero point for goodness entails that an empty life lived in an entirely unconscious state is neither good nor bad – it is neutral. If Parfit is right, asking whether the never-conscious life is a bad life should sound to us like the absurdity, "Is a life that is neither good nor bad a bad life?" Yet Broome thinks it is possible to conclude that such an empty life is bad rather than neutral. At the very least, it is an intelligible question whether the life is bad. Therefore, we must reject Parfit's zero point for goodness which deems the question absurd.

One response we might make is that there can be a zero point for good and bad without this being a hedonistic zero point[5]. There can be a zero point for good and bad on an objective list theory of wellbeing, but an empty life could still qualify as a bad one if long periods of unconsciousness are one of the bads on the objective list.

Even if we favor a hedonistic zero point for value, Broome's use of Moore's open question argument is not compelling.

In Broome's hands, the open question argument assumes that a particular question that does not initially look like a tautology ought therefore to be considered an open question, and so rules out any suggestion that would answer and thus close and tautologize the question. This means we can use the open question argument against any proposal that gives what we think to be the wrong answer to a question, or which gives any answer at all to a question we are uncertain about.

Any question can seem like a tautology if we think we know the answer. Imagine someone asks, "Is The University of Oxford the best institute of higher learning in Oxfordshire?" If we think the answer is yes, this question could sound like asking, "Is The University of Oxford The University of Oxford?" or "Is the best institute of higher learning in Oxfordshire the best institute of higher learning in Oxfordshire?" If, however, we believe this is an open question – which we are more likely to do if we are uncertain of the answer or if we think Oxford Brookes is better than The University of Oxford – then the question will not sound like a tautology, and

[5] This point was suggested by [redacted].

we will reject any proposal that suggests it should[6].

Moore proposed an extensive list of general goods in place of the classical utilitarian mono proposal of pleasure. If we asked, "Are [every general item of good on Moore's list] all the good things?" this would sound to Moore like "Are all the good things all the good things?" However, "Are [every general item of good on Moore's list] all the good things?" will sound to the unconverted like an open question. And thus, if we accept the open question argument, we therefore must reject Moore's proposed list of general goods that turns this question into an absurdity.

The open question argument can also be employed against Broome. We might wonder, "Can something be better without being good, and can something be worse without being bad?" For Broome this sounds like asking, "Can something be better without being better, and can something be worse without being worse?" Given Broome's assumptions, this becomes an absurd question. But "Can something be better without being good, and can something be worse without being bad" is not obviously an absurd question. Similarly, "Does good reduce to better?" sounds to Broome like, "Does better reduce to better?" or "Does good reduce to good?" But "does good reduce to better?" seems like an open question – it is at least an intelligible one. And so if we accept the open question argument, we must reject Broome's proposal.

Most questions like these do not actually sound like tautologies, even to those who believe in definitive answers to these questions. This does not prove that Moore

[6] Even if we think Oxford Brookes is the best institute of higher learning in Oxfordshire, the question would sound absurd if we make our desired replacements in the question. An Oxford Brookes devotee would hear, "Is The University of Oxford Oxford Brookes?"

and Broome are right that these are therefore open questions and that we must reject proposals on the grounds that they suggest otherwise. Questions that we believe to be closed can nevertheless sound like intelligible questions because we recognize that there is disagreement and uncertainty about what the answers to these questions are.

I therefore second Dancy's rejection of the open question argument. Yet rejecting the open question argument creates a minor dilemma for me. Employing the open question argument against Broome's own claim, as I just have, reveals that part of my argument against Broome's value theory has been relying on something like the open question argument too. Recall the person in the pain/pleasure chair whose distinctions between bad and worse were rendered into absurd redundancies by Broome's reduction of all value into comparative value. In presenting this as a problem for Broome, I was assuming there is a distinction between good and better and between bad and worse, which Broome's reduction of good to better and bad to worse obviously fails to capture. Those who think there is a distinction between good and better and between bad and worse will agree with me that "That's worse. In fact, this is downright bad" is a sensible sentence. But those who agree with Broome that all value is comparative will side with Broome that this is a confused thing to say.

Nevertheless, Broome has provided no compelling argument for the reduceability of good and bad to better and worse. The goods and bads that we seem to bestow by creating new lives, and which are not better or worse for those who come into existence, stand as plausible refutations of the claim. Broome gives us no reason to think that this question is a closed one.

2.2 Non-Comparative Value is Action Guiding

Someone might accept that good and bad do not reduce to their comparatives, but insist that non-comparative value is not action guiding, and it is only comparative value that matters. McMahan has called this the "comparative view" (1981:104). Narveson (2012[1978]) may be one proponent of the comparative view. I will attempt to refute the comparative view in this section.

Why should we care about non-comparative value? Bykvist's answer is, "In general, in a choice between a good x and a y that lacks any value (is neither good, bad, nor neutral), it is reasonable to prefer x to y. So, you can care about goodness and prefer x to y without x being better than y" (2015:93).

I would add that better and worse cannot matter if their constituent components, good and bad, do not. If someone's good did not matter, more or less of it could not matter, and if someone's bad did not matter, more or less of it could not matter. Zero plus or times zero is still zero. If better and worse matter, then good and bad must as well, because better and worse refer to levels of good and bad.

This might just be entailed by how I have defined good, bad, better and worse. Different definitions could make this a trickier matter. However, aside from good and bad reducing to their comparatives, I am not sure what alternative definitions we could devise for these evaluative terms that would rule out this sort of response.

My claim, then is this: if better and worse are action guiding, then non-comparative good and bad must be action guiding as well. This means that the good of a potential life is a reason for creating a life, and the bad of a potential life is a reason

against creating the life. How might this be denied?

Someone might say that good and bad matter, but only when these constitute better and worse; when good and bad come apart from better and worse, as these do with regard to creating new lives, then good and bad do not matter anymore. But this looks like an *ad hoc* attempt to deny that there can be value in creation. Why should we think that states themselves do not matter on their own, but that moving between states does? If better and worse refer to amounts of good and bad, why should a change in amounts of good and bad be action guiding while good and bad themselves are not?

If non-comparative value were action guiding, it would help to explain some of our evaluations about creating lives, extending lives, and shortening lives. Many people think it is wrong to create a miserable life. Anyone who is concerned with personal value rather than impersonal value, and who does not think we can compare existence and nonexistence in personal terms, and who thinks non-comparative value either does not exist or does not matter, will have to give a non-axiological explanation for this wrongness. Bader says creating a bad life violates a consistency constraint about the reasons we have, whereas others say that creating a bad life is a rights violation, or makes things impersonally worse, or is simply wrong. Yet there seems to be a personal-value-based reason for not creating a bad life. Creating a bad life does not seem like an impersonal affront to the universe. It is wrong because this life will be bad for someone. When someone with a good life wants to live longer, they are thinking that this extra life will be good for them; they are not worried that their life ending prematurely will be impersonally bad. When someone has a horrible life and wants to shorten it, they are

not thinking that ending their life will make things impersonally better; they are thinking that they can end the badness of their lives if they cease to exist.

Non-comparative value seems to be action guiding. In the next section, I will say more about just how action guiding it is, particularly with reference to comparative value.

2.3 Choosing Between Comparative and Non-Comparative Value

That non-comparative value exists and is action guiding is essential to the Epicurean Reconciliation Strategy (ERS), but there is still a question of how to weigh non-comparative value against comparative value. If good and bad are not reduceable to their comparatives of better and worse, then these may sometimes come apart. This raises the question of whether it is good and bad or better and worse that matter most when value splits and we have to decide whether to prioritize comparative or non-comparative value. Is there any reason to think that some amount of good or bad that is also better or worse is more important than good and bad of the same magnitude that is not better or worse? McMahan sees this as an important question that arises when we can choose whether to extend an existing life or create a new life (2013). He treats the extension of an existing life as comparatively better or worse (as well as non-comparatively good and bad), while treating the creation of a new life as only non-comparatively good and bad. He believes the comparative benefit of extending an already existing life is the more essential, and that this could be because it is a comparative benefit rather than a purely non-comparative one. This suggests that better

and worse take priority over good and bad if we have to choose.

I claim that extending a life provides only non-comparative benefits just as creating a new life does. Therefore, to express this question we are considering in a more universal manner, we could consider a conflict between adding a certain amount of pleasure to an already existing life – which we all agree would be a comparative benefit – and creating a new life with that same amount of pleasure instead (which I claim is a non-comparative benefit). Many people have the intuition that comparative benefits are the more valuable of the two – that we ought to add the pleasure to the already existing life rather than create a new life with that same amount of pleasure. In this section, I will argue that when our concerns are purely axiological, it is only the magnitude of non-comparative good and bad that matter. Comparative value matters derivatively, to the extent that non-comparative value is at stake. Non-comparative value is always action guiding. Comparative value that is empty of non-comparative value could be irrelevant.

I will defend this claim by appealing to what I see as the axiological grounding of Parfit's "No Difference View." Parfit introduced the No Difference View (NDV) in (1984) and elaborated on it in (2011:6-7). He treats it as a fairly narrow judgment of a particular kind of case, and he does not seem interested in extending the view beyond this context. After summarizing how Parfit uses the NDV, I will describe and defend what I see as its axiological grounding. Once identified, the Axiological Grounding of the No Difference View (AGNDV) can be applied more broadly than the NDV that it entails. Specifically, I will use AGNDV – and my defense of it – to show that

comparative value does not take priority over non-comparative value, and that this has wide-ranging implications, some of which are non-intuitive.

Before I can describe what I see as the AGNDV, we need to consider the NDV itself. Parfit discusses NDV in the context of non-identity cases. These are scenarios in which our actions affect more than the value of outcomes – our actions also potentially affect the identities of those who exist to receive the value in the outcomes, because our actions alter who will be born. A general sort of case we can imagine in which our actions influence identity in this way is a case in which we can bring about a certain amount of value because we improve someone's circumstances, lifting them to a certain level of wellbeing (six, say) that they would not be at otherwise, or we can bring about that same amount of value because we create someone so that they are in these same circumstances that provide that wellbeing level of six. The first disjunction involves comparative value; it improves someone's circumstances – it makes them better off. The second disjunction involves only non-comparative value. It creates wellbeing that is not better for anyone, but which is just as good as the wellbeing that would be better for someone if someone who already existed at a lower level of wellbeing were brought to this level. Parfit's No Difference View is simply this: in "same number cases" – that is, cases in which the same number of people will exist no matter which action we choose – it is the value of the outcome that matters, and not the identities of those who receive the value, and not whether the value will be better or worse for them rather than just non-comparatively good or bad for them.

In (1984), Parfit applies NDV to a case he calls The Medical Programmes. In it,

we must decide between cancelling a medical program that will prevent 1,000 handicaps by curing pregnant mothers of a condition that would otherwise handicap their already developing children ("Pregnancy Testing"), and a medical program which will prevent the same number of the same type of handicap by encouraging 1,000 women to delay pregnancy for two months – as this is how long it takes for a disease they have to disappear – thus conceiving different non-handicapped children than the handicapped children they would have conceived during those two months in which they had the condition ("Pre-Conception Testing"). If we cancel Pregnancy Testing, this will be worse for the 1,000 already-conceived children who will be born with a handicap that could have been avoided. It will not be better for anyone because the children who are conceived free of handicaps due to Pre-Conception Testing would simply not have existed if Pre-Conception Testing had been canceled instead; their being conceived without a handicap is not an improvement of their situation, it is the creation of their situation. If we instead cancel Pre-Conception Testing, this will be better for the already-conceived children who will therefore be prevented from developing handicaps (thanks to Pregnancy Testing), and it will be worse for no one, because those who are born with the handicaps could not have existed free of these handicaps; they would never have existed at all if their mothers had waited two months and conceived other children. No matter which program we choose to cancel, 2,000 children will be born, and half of them will have this particular handicap.

Parfit gives his NDV-inspired verdict:

> I therefore ask: 'If there will be people with some handicap, the fact that they are handicapped is bad. Would it be worse if, unknown to them, their handicap could have been cured?' This would be worse if this fact made these people worse off than the people whose handicap could not have been cured. But this fact does not have this effect. … Since this is so, I judge these two outcomes to be morally equivalent. Given the details of the case, it seems to me irrelevant that one of the groups but not the other could have been cured (1984:368).

Having a particular handicap when you would not have had the handicap if things had gone differently for you (and you do not know you would not have had it if things had gone differently for you) is not worse than having that same handicap when that handicap was inevitable for you (and you do not know it was inevitable for you). The worseness of something bad for someone does not make that bad thing any more disvaluable for them than the same bad thing would be for someone for whom that bad thing is not worse. The betterness of something good for someone does not make that good thing any more valuable for them than the same good thing would be for someone for whom that good thing is not better. Unknown counterfactual alternatives to something good or bad do not affect that thing's value. Better or worse nearby possible worlds that are also undiscovered worlds do not affect the value of this world. The value of my life is determined by how my life goes in this world. All the possible ways in which things could have been different do not influence the value of the life I actually have – assuming, again, that I do not know about these foreclosed possibilities. What

would be better for us in the future and what would be worse for us in the future are certainly worth considering, because these possibilities can influence our decisions and help us achieve a life with more intrinsic or non-comparative value. Thinking about better and worse prospectively can help us to actualize a possible world with less bad and more good. But what *would have been better* and *would have been worse*, or what *is better now* and *is worse now* is irrelevant to assessing the value of the life that actually did occur in defiance of other possibilities, if we think that nearby possible worlds do not leak value into this one.

That, I take it, is the Axiological Grounding of the No Difference View (AGNDV). I will now attempt to show that we should accept AGNDV, and therefore accept that comparative value (better and worse) is not axiologically more important than non-comparative value (good and bad). Consider these hypothetical scenarios:

A must press one of three buttons. A presses the middle button, and this puts A in a state of level-five pleasure for an hour. Both other buttons would have put A in a state of level-nine pleasure for an hour. Pressing the middle button was good for A but was also worse for A. A never finds out what the other buttons would have done.

B must press one of three buttons. B presses the middle button, and this puts B in a state of level-five pleasure for an hour. Both other buttons would have put B in a state of mild pain for an hour. Pressing the middle button was good for B and was also better for B. B never finds out what the other buttons would have done.

C must press one of three buttons. C presses the middle button, and this puts C in a state of level-five pleasure for an hour. Both other buttons would also have put C in

a state of level-five pleasure for an hour. Pressing the middle button was good for C but was neither better nor worse for C. C never finds out what the other buttons would have done.

AGNDV tells us that each of these scenarios has the same value, despite the variations in betterness and worseness. Prudentially speaking, we should be indifferent between being A and pressing the middle button, being B and pressing the middle button and being C and pressing the middle button. Each of them feels level-five pleasure for an hour and never finds out about the alternatives that were open to them. Whether this same level of pleasure can be described as better or worse for them is prudentially irrelevant.

AGNDV does not entail a commitment to a non-comparative account of harm. Gardner writes that a non-comparative account of harm holds that, "harm is a state of affairs that is bad for an individual, such that the badness of the state does not derive from its being worse for the individual than some other state that the individual was or would have been in" (2016:430). AGNDV is consistent with the more widely held counterfactual account of harm in which something may harm someone by making them better or worse off than they would otherwise have been. If a counterfactual account of harm is accepted, AGNDV will entail that something that is bad for someone which is arrived at through this individual's being counterfactually harmed is not more disvaluable for this individual than the same bad thing would be for someone who arrived at this bad thing without their being counterfactually harmed. This is assuming unawareness. If I have a handicap because someone counterfactually harmed me, and I

do not know this harm is why I am handicapped, my being handicapped is not more disvaluable for me than the same handicap is for someone else who got that handicap without being counterfactually harmed. It is intrinsic features from the perspective of the harmed or not harmed agent, and not counterfactual alternatives, that determine value in AGNDV.

AGNDV is consistent with the position that imposing something bad in a counterfactually harmful way is impermissible while imposing that same bad in a way that is not counterfactually harmful is permitted. However, the No Difference View (NDV) may not be consistent with this deontological agnosticism. It might be that NDV would have us translate AGNDV directly into the rightness or wrongness of actions in same-number non-identity cases. If so, there would be no moral difference between imposing a bad on someone in a counterfactually harmful way and imposing that same bad on someone without this being counterfactually harmful. The advocate of AGNDV does not need to endorse this stance. It is nevertheless worth exploring the stance, to see what might be appealing or unappealing about it, and to more clearly distinguish the AGNDV from it.

Weinberg seems to interpret NDC as implying the stance I just described (2014). Weinberg attacks NDV by having us consider different scenarios in which a lesbian couple ensures that their child is deaf. In one scenario, they secure the deafness of their child by getting sperm from a deaf donor, and they name the resulting child Conan. As Conan's deafness was genetically determined, Conan could not have existed without being deaf, and so coming into existence with a soundless life worth living is not worse

for Conan. In the contrasting scenario, the couple conceives a non-deaf child named Indy, and they make Indy deaf by giving him a deafness-inducing medicine before he ever hears anything.[7] If they had not given Indy the medicine, he could have spent his whole life hearing. If we assume it is generally worse to be deaf than to have hearing, causing Indy to be deaf makes him worse off than he otherwise could have been. This is not the case for Conan, because being deaf is a necessary part of Conan's identity and existence. To rule out some possibly confounding factors, like a worry that Conan's parents may be better at parenting than Indy's parents are, imagine that as soon as Conan is born deaf due to his genes, his parents change their mind about raising a child and give him up for adoption to Family A. Imagine as well that in the contrasting scenario, in which it is Indy who is born and then is immediately made deaf through a medicine, his parents change their mind about raising a child and give him up for adoption to Family B, which is relevantly similar to Family A.

Weinberg takes NDV to entail that purposely combining genes in such a way that a child is born deaf is the same, morally, as conceiving a hearing child and then painlessly making that child deaf – a result that violates Weinberg's intuitions. AGNDV is axiological rather than deontic, and is therefore agnostic about whether there is a moral difference between intentionally conceiving an inevitably deaf child and conceiving a hearing child and instantly and painlessly making that child deaf. Of Conan and Indy, AGNDV simply says this: it is no worse to be Indy than it is to be Conan, even though Conan's deafness was built into him (and thus is not

[7] Weinberg does not specify that Indy goes deaf before hearing anything, but I take it Weinberg had something like that in mind.

counterfactually worse for him) and Indy's deafness was externally imposed on him (and thus is counterfactually worse for him). The claim is that any disvalue of deafness consists in how deafness affects one's life in this world. Whether or not there is a nearby possible world in which one can hear does not affect one's life in this world. It may still be the case that the deontic verdicts of the two cases differ.

I believe the AGNDV arrives at plausible verdicts, and that we should accept it. We have no axiological reason to prioritize better and worse over good and bad. If this sounds counter-intuitive, this may in part be because "better" often refers to the most good-producing action available, and an instance of just plain "good" might not usually amount to as much good as an instance of "better" might. We should not let this typical association confuse us. Something that is non-comparatively good may involve more goodness for beings than something that can be called better. Our heuristic preference for better over good might not always lead to the most favorable outcomes.

The AGNDV suggests that we have reasons to do what is good and reasons to avoid doing what is bad, but our reasons to do what is better or worse are dependent on how this betterness or worseness relates to non-comparative goodness or badness. If something is better because it is more good for beings or less bad for beings compared to an available alternative, then we have reason to do it. If something is worse because it decreases good for beings or increases bad for beings compared to an available alternative, then we have reason not to do it. If something is better for someone but it neither increases good for beings nor decreases bad for beings compared to an available alternative, a coin flip is in order. This is true as well if something is worse for someone

but it neither decreases good for someone nor increases bad for someone compared to an available alternative. However, if something is not better for someone, but it is good for someone, we have reason to do it, and if something is not worse for someone but it is bad for someone, then we have reason to avoid doing it. Creating a miserable life is a salient example of this latter point. It is not worse for the being who comes into existence that they exist, but it is bad for them, and that badness is the reason to not create them. If it were also somehow worse for this miserable being to exist, that worseness would add no axiological reasons to avoid the creation that the badness itself did not already provide.

This prioritizing of intrinsic rather than comparative value inspires my revision to the standard phrasing of the person-affecting view that something cannot be better or worse if it is not better or worse for someone (Temkin 2000:132). Instead we say that something cannot be good or bad if it is not good or bad for someone.

In the next section, I will combine the ideas I have laid out so far to form my version of the Epicurean Reconciliation Strategy (ERS).

2.4 The Epicurean Reconciliation Strategy

In the first chapter, I established that coming into existence and continuing to exist are neither better nor worse for those who exist, and that death is neither good, bad, better nor worse for those who cease to exist. Combined with a person-affecting view, that it is personal goodness and badness that matter morally, this may seem to entail full neutrality about existence: that there is no value in coming into existence or in

continuing to exist. However, I argued earlier in this chapter that there is non-comparative value of good and bad, and that this value is action guiding. This allows me to say that there is no comparative value in coming into existence and continuing to exist, but coming into existence and staying in existence are non-comparatively good or bad for those who do so, and this non-comparative value gives us reasons to create or to avoid creating lives, and reasons to extend or to shorten lives.

When McMahan introduced a view like this, he called it the Epicurean Reconciliation Strategy (ERS). McMahan described the core ideas behind it without endorsing it himself:

> [I]f we bring a person into existence with a life that is worth living, there is no problem in locating the beneficiary. He is here before us, enjoying the goods of life. There is thus no obvious incompatibility between the claim that it can be good for a person to come into existence and the [Epicurean] Existence Requirement [that someone must exist for anything to be good, bad, better or worse for them]. There is, moreover, a relevant parallel here between starting to exist and continuing to exist. You might have died five years ago. But you did not; instead you continued to exist. If your life has been worth living, then it was good for you to have continued to exist. There is no problem in this case in locating the subject of this good fortune: you are the beneficiary of your continued existence, and you clearly exist during the time when this good fortune occurs. Hence it is not obviously incompatible with the Existence

Requirement to claim that your continuing to exist is a good thing for you.

> The claim of the reconciliation strategy is that continuing to live can be good for a person even if, as the Existence Requirement implies, death would not be bad for him. … This…allows us to concede the Epicurean claim that death is not bad for people, or worse for them than continuing to live, while at the same time accepting that continuing to live can be good for them. We can say that continuing to live can be a good thing for a person, though it is not better than dying. In short, we simply give up the comparative claims, while retaining the noncomparative ones. (1988:35, 37)

So far this thesis has argued for the premises that I see as leading to ERS. Even if we accept these premises, we might not like ERS itself. I will now consider some implications of ERS that may initially seem unintuitive enough that they may tempt us to abandon ERS. I will try to show that these implications are not as bad as they seem, and that we should not abandon ERS because of them.

A possible intuitive problem for ERS is that it does not allow us to talk about better or worse lives in some cases in which this comparative language seems appropriate. This is one of the results of ERS acceptance of what Bader calls the temporal part affecting view (TPAV), and which I presented as the view that longer and shorter lives cannot be fully compared without implying an impossible comparison between existence and nonexistence. Consider two lives of the exact same length. Life 1

is barely worth living because it has a constant but meager amount of good in it, and no bad in it. Life 2 is very much worth living because it has a constant and large amount of good in it, and no bad in it. ERS uncontroversially entails that any given moment of Life 2 is better than any given moment of Life 1, and any given moment of Life 1 is worse than any given moment in Life 2. Furthermore, Life 2 is overall better than Life 1, and Life 1 is overall worse than Life 2. Here, ERS agrees with most value theories. Now imagine Life 3, which has the exact same amount of good in it as Life 2 does, but Life 3's goods are less concentrated; they are spread over a longer lifespan so that any particular moment of Life 3 is as barely worth living as any particular moment of Life 1. Any segment from Life 3 that is no longer than Life 1 is overall comparable as equally valuable as a segment of equivalent length from Life 1. But the whole of Life 3 cannot be compared with Life 1, because Life 1 is not long enough for a full overall comparison. Even though Life 3 has the same amount of good in it as Life 2 does, and this is more than the amount of good in Life 1, we cannot say that Life 3 is better than Life 1 – it is not better at any particular point, and endorsing something like TPAV means we cannot make a comparison that would deem it better overall. Life 3 is in part equal in value to Life 1, and it is in part incomparable to Life 1, despite Life 3 having more value in it than Life 1 does.

The first problem with this is that we might think the comparison between Life 3 and Life 1 should match the comparison between Life 2 and Life 1, and yet ERS does not allow this. The second problem is that it may also seem that we should favor the creation of Life 3 over Life 1, but that if we cannot call Life 3 better than Life 1,

perhaps we cannot do this, even though we can favor the creation of Life 2 over Life 1.

The second problem is not worrisome for ERS, given that ERS holds that non-comparative value is action guiding, and indeed that non-comparative value is potentially more significant than comparative value is. Accepting ERS means we can call Life 3 "more non-comparatively good" than Life 1, or just "more good" than Life 1, and favor creating Life 3 over Life 1 because of Life 3's greater goodness.

Regarding the first problem, the Reconciled Epicurean can say that perhaps it is appropriate that the Life 1/Life 2 and Life 1/Life 3 comparisons do not match, given that Life 2 and Life 3 are different sorts of lives despite having the same amount of good in them. It is rational for someone to prefer to live Life 3 over Life 2 and vice versa, but it is not rational for someone to prefer to live Life 4 over Life 5 if Life 4 and Life 5 are identical in every respect.

In fact, that Life 3 and Life 2 do not compare with Life 1 in the same way, and do not fully compare with each other, may be an intuitive advantage for ERS. Most individuals are not indifferent between lives of different lengths with the same amount of goods in them in different concentrations. ERS here reflects our ambivalence and confusion about how to compare lives of different lengths and different patterns and concentrations of goods and bads[8].

Still, there might be some kind of problem here. ERS does not on its own give us the resources to choose between creating Life 2 and Life 3. If we have a strong

[8] Temkin thoroughly explores this kind of ambivalence and confusion in (2012), and which I do not have the space to explore more here

preference toward one or the other, we might want ERS to reflect that preference, but instead it leaves the choice up to us. If we can only create one life, and we know Life 2 and Life 3 are equal in non-comparative goodness and non-comparative badness, ERS does not pick a side. On its own, ERS would even tell us to pick Life 6 over Life 7 if Life 6 is a long life that is barely worth living but contains 100 good, and Life 7 is a short and blissful life that is better than Life 6 at every comparable moment but contains only 99 good. There is some linguistic oddness here as well, in that ERS entails that Life 6 and Life 7 are not comparable overall, and that Life 7 is better than Life 6 to the extent that these lives are comparable – and yet that we should favor the creation of Life 6. The Reconciled Epicurean who wants to avoid this intrapersonal repugnant conclusion will need to bring in other considerations having to do with how we aggregate wellbeing within and across lives. I do not have the space to discuss those possibilities now.

One reason for accepting ERS despite any intuitively problematic implications it might have is that it supports the plausible view that death is not worse than never existing, and death is therefore not in itself a reason to avoid creating a new life – even when the lived life is neutral and the death prevents what would have been a very good life. Kamm defends this intuition without accepting ERS (1992). Like Nagel (1970), Kamm says that death can be bad and worse for those who die, and that this badness is purely comparative, because there is nothing intrinsically bad about death. This lack of intrinsic badness is why death is not a reason to avoid creating life. She writes:

> Death is bad for a person primarily because it deprives him of more life. If this is true, then why should it be better for him, from the point of view of his best interests, to have had no life at all than to have had a short life? It is true that we are susceptible to the harm of being deprived of more life only because we do have some life. Yet, the suggestion is, avoiding this harm of deprivation does not seem to be a reason to avoid creating some life. That is, in this case one is not harmed overall by coming to exist and then having the harm of being deprived of life happen. … In short, death, including its deprivation of more life, may be the sort of bad thing that it makes no sense to avoid for the sake of the [being] created, by not creating it at all. (1992: 84-85)

While Kamm shows that an anti-Epicurean can maintain that death is not a reason to not create a life, ERS is linguistically superior at conveying this, and also offers a clearer substantive explanation for why death is not a reason to not create life. Consider the case of a child named A who is born into a life that is neutral and remains neutral for as long as she is alive. She dies at age 5, and it just so happens that if she had kept living, her life would have gone from neutral to fantastic for the next 95 years. If we accept Kamm's anti-Epicurean view and terminology, we would say that A's life was completely neutral for her, and that death was very bad and worse for her. Since the only evaluations we can make regarding A are neutrality badness, and worseness, it seems we should be opposed to creating A. And yet, Kamm seemingly says (and ERS would have us agree) that we should be axiologically neutral about creating A. This

badness of death terminology that Kamm is committed to ends up being confusing when she combines it with the claim that death is not a reason to not create life. ERS does better here. It would have us say that A's life is neutral, and that A's death is valueless. This makes it clear that we should be axiologically neutral about creating A. ERS is also clear about exactly why it is that death is not a reason to avoid creating life. For the Reconciled Epicurean, death is not bad for the one who dies, and it is not worse for the one who dies. Death can be less good for the one who dies, but this less goodness is a non-experienced absence of good which cannot be distinguished from the absence of good in never existing. Reconciled Epicureans view existing sentient beings as essentially "resettable" to nonexistence. Someone's death recreates the conditions of that individual never having existed.

However, the ERS explanation for the intuition that death is not a reason to avoid creating life – combined with the ERS platform that non-comparative good is the value that matters, and creating lives and extending lives provides the same sort of non-comparative good – means Reconciled Epicureans are *prima facie* unable to favor the currently existent over the potentially existent, if we are ignoring negative side effects of untimely deaths on existing beings. This means that ERS entails the "replaceability" of all beings, including rational ones. In the next chapter, I will explain what replaceability is, why it seems objectionable, why replaceability may actually be axiologically desirable, and how ERS can avoid replaceability in deontic practice.

Chapter 3. Defending the ERS from the Replaceability Objection

Impersonalist total hedonic utilitarianism seems to obligate the destruction of existing sentient beings in favor of creating new sentient beings whose lives are better than the originals' lives, whenever this would maximize aggregate utility. An objection to ERS is that despite being a wide person-affecting theory, it seems to obligate the same thing. In 3.1, I will define replaceability and explain why ERS – as a purely axiological theory – is indeed committed to it. In 3.2, I will consider and reject proposals from Dworkin and Cohen which would provide an axiological barrier against replaceability by placing greater value on what already exists. In 3.3, I argue that any axiological opposition to replaceability will have implausible implications. If we are to avoid replaceability, we must appeal to non-axiological considerations. In 3.4, I will consider a non-axiological anti-replaceability proposal from Velleman (2015) and others which holds that value matters only if the value bearers already exist and matter. Though I will reject this idea, it will inform my own ERS-compatible solution to replaceability. In 3.5, I will propose my own deontological complement to ERS that gives us moral obligations to the currently existent but not to the potentially existent, allowing us to avoid replaceability without abandoning ERS as an axiological theory.

3.1 What Replaceability Means, and Why ERS Entails it

Moral theories are said to permit or obligate replaceability when they hold that the value

of someone's life for that being is replaceable by the creation of a new being whose life is equally valuable or more valuable to the new being (Uniacke:1997). In impersonalist total hedonic utilitarianism, the feelings of sentient beings serve to increase the aggregate of good feelings. If what matters morally is improving or maintaining a certain aggregate wellbeing score, the replacement of one happy being with another equally happy being is axiologically neutral so long as the aggregate wellbeing score never budges. The role of adding to the overall goodness in the universe has been taken over by the new being at the moment of its creation, and so we should be indifferent between the continuation of a being or the painless replacement of that being with another being of equal value (Singer 2011:106). This is one reason many people reject impersonalist total hedonistic utilitarianism.

By focusing on personal value rather than impersonal value, Reconciled Epicureans avoid instrumentalizing the experiences of sentient beings. When we destroy one life to create a new life with the same value, we replace what is valuable for one being with what is valuable for another being, but this does not compensate the first being for the destruction of its good. Since value for beings is the only moral value that strictly person-affecting views like ERS recognize, this could seem to make replacement impermissible. However, ERS is not a narrow person-affecting view; it recognizes non-comparative value, and it does not distinguish between future value for merely possible beings and for already existing beings. ERS thereby avoids indifference to death, creation and extinction, but these same features mean ERS looks committed to

the Repugnant Conclusion[9] and replaceability.

If we accept that goodness that is better for someone is no more valuable than goodness that is not better for someone, and badness that is worse for someone is no more disvaluable than badness that is not worse for someone, then previously having had consciousness is irrelevant to determining the value of future consciousness. Just as ERS leaves no room for a prudentially relevant difference between nonexistence for the never born and nonexistence for the once existent, the Reconciled Epicurean does not see any prudentially relevant difference between prospective future life for an existing being and prospective future life for someone we might create. The timing of a life and its degree of contingency – whether its existence is dependent on or independent of our actions, for instance – are irrelevant to the life's non-comparative value. Future existence would be good for an existing individual, but future existence would also be good in the same way for a new individual we can create. And future existence for an existing individual would not be better for them than their ceasing to exist, either because of the symmetry of the missing relatum, or because if they ceased to exist there would be no one for whom this lack of good would be worse. Even if future existence would be better for the currently existent and not better for the possibly existent, this would not matter given that my version of ERS denies that comparative value is more important than non-comparative value.

Imagine X exists and Y does not exist but could be made to exist. It would be good for X to be alive tomorrow. It would be good for hypothetical non-existent

[9] I will not discuss the Repugnant Conclusion in this thesis.

individual Y to be alive tomorrow. It would not be good for X if X did not exist tomorrow, and it would not be good for hypothetical non-existent individual Y if Y did not exist tomorrow. If X ceases to exist and Y does not come into existence, there is no difference in this non-goodness for formerly existent X and never existent Y; ERS entails that X would not be missing out on tomorrow in a different way than Y would be. Both had potential to be alive tomorrow and both do not exist tomorrow and this is why there is no good of tomorrow for X or Y. The absence-of-tomorrow blank is the same whether or not there was a yesterday or today. For the Epicurean Reconciler, the only difference between these two possible good tomorrows or lack of tomorrows is that with X, we already have a face to put on the individual who will or will not go on to experience the good[10].

The value of creating and extending lives is axiologically interchangeable when we do not think that someone's having existed now and in the past tells us anything about the future value of that life that distinguishes it from the value of an equally good new life that we could create. If the presumed good future life of an existing being counts in favor of extending that life, the presumed good future life of a potential being counts in favor of creating that life – and to the same extent, if we are not thinking about effects on others. While previous life can add to the richness of one's future life, and a newly created life must go through a lengthy development process – and so a life that has just been created is qualitatively different from a life that is extended – a new

[10] This may partially be why we intuitively think more is at stake in X's good tomorrow than in Y's good tomorrow.

being we can create will presumably have a richer future life as their life continues. Already existing beings are not unique in this way; they just happen to be further along in the process of achieving greater value.

Replaceability seems entailed by ERS whenever a merely possible life has more good in its prospective future than an existing life does, and when the creation of the possible life depends on the destruction of the existing life. Uniacke writes, "Proponents of…replaceability argue that if we are justified in acting for improvement in not creating some possible individuals in favour of creating others, then we are justified in acting on this basis in destroying some actual individuals in favour of creating others. They deny the intuitive moral distinction between these two courses of action…" (1997:161-162). If we replace "moral" with "axiological" in that last sentence, this perfectly describes the Reconciled Epicurean, because the unadorned ERS treats the destruction of a good life and substitution with an even more good life as axiologically similar to creating the more good life rather than the less good life when neither exists yet[11]. As far as the Reconciled Epicurean is concerned, regarding creation and extension decisions, it may as well be that no one exists yet. (Again, assuming we are ignoring side-effects, and we are thinking axiologically.)

As we have seen, this is because the destruction of the first being is not bad or worse for that being, and the benefit for the second being of coming into existence is not necessarily inferior in value compared to the value for the first being of continuing

[11] The complication with this kind of example is that if we do a replacement, we have to count the value of the lived portion of the replaced life, along with the value of the replacing life, which leads to a different amount of value than if we had created the more good life initially and made no replacement.

its own life. The unique individuality of the first being is not more valuable than the unique individuality of the possible second being just because one came into existence first. The first life is irreplaceable in the sense that it is unique and no new being we create will truly take the first life's place, and yet we can have more value-based reason to make possible the future existence of a previously non-existent life than to allow the continuation of the already existent life, because the merely possible life has better prospects. What distinguishes ERS from the impersonalist total view axiologically is that this replaceability is entailed by non-instrumental reasons of promoting value for beings.

Imagine there is a planet with one happy rational being on it. Pressing a button painlessly kills this being and replaces her with a new being whose future prospects are more good. If we press this button, we will have done what is bad for no one, and what is worse for no one. We have destroyed good for the first being, but we have also created good for the new being, and this good for the new being is greater than the good the first being would have had if we had not pressed the button. The unadorned ERS has us press this button.

The ERS version of replaceability parallels lifeboat cases in which not every life can be preserved, and some of the lifeboat occupants have better prospects than others. Preserving the good future life of the teenager who is allowed to stay onboard does not help the older person with a chronic disease who is thrown overboard. The value of the destroyed life is not truly replaced by the extension of the extended life. Yet killing one to extend the other can be the best thing to do. If superior prospects give us reason to

extend one existing life rather than another existing life, and superior prospects give us reason to turn one potential life actual rather than turn another potential life actual, then for the Reconciled Epicurean, superior prospects can also give us reason to create a new life instead of extending an existing life[12]. Extending a life is just making possible some future life, and creating a new individual is also just making possible some future life.

The lifeboat analogy offers another angle on Parfit's non-identity problem. Just as throwing the teenager overboard instead of throwing the sickly older person overboard is the wrong choice even though it is good for the sickly older person for her life to be saved if her life is worth living,[13] creating a less happy life that is worth living over a happier life that is more worth living violates our axiological reasons even though the less happy life benefits from being created. There is nothing puzzling about the wrongness of saving the sickly older person instead of the teenager despite this saving being good for the sickly older person, and for the Reconciled Epicurean, this directly carries over to how we should think about the axiology of creating or extending lives. In cases of extension or creation, the bare ERS tells us to choose whichever leads to the most non-comparative good.

If the lifeboat analogy makes replaceability seem a little less alien compared to the rest of our views, and therefore less disturbing, the Reconciled Epicurean might want to go further and argue that we should not trust our anti-replaceability intuitions at

[12] However, there are also plenty of cases in which the Reconciled Epicurean thinks we ought to save the life of an adult rather than create a new life, since creation is costly, babies are unable to take care of themselves, and the potential creators may be unable or unwilling to care for the baby for what could end up being the rest of their lives.

[13] And, the Epicurean would add, it is not bad or worse for the teenager who is thrown overboard to die, so long as she dies instantly.

all. One way to do this might be to attribute the intuitions to morally irrelevant psychological features of ourselves. For instance, a factor that might be corrupting our intuitions is that benefitting existing individuals may feel more similar to helping nearby individuals we know, whereas benefitting individuals by bringing them into existence feels a bit more like benefitting distant strangers. This is important because benefitting nearby individuals tends to be more emotionally salient; the same may be true for benefitting existing beings, and this could partially explain why our intuitions revolt against the axiological equivalence of creating and extending lives when the same amount of good is at stake.

Another reason we tend to think it is more important to help existing beings instead of creating new beings may be that we think existence is generally difficult and that existing beings are in need of a lot attention and resources for their lives to avoid becoming bad, and so we expect a policy of creating new lives instead of helping the existent to lead to mediocre lives for the existent. This could relate to the plausible and widespread asymmetric intuition about good and bad – that it is more important to prevent bad things than to promote good things. Someone who does not exist cannot be deprived of wellbeing improvements, and cannot suffer drops in wellbeing. There is no suffering involved in not creating someone, but there could be suffering for existing beings if we are too focused on making new individuals to care for the individuals who already exist. This could contribute to our general distaste for replaceability even though we can imagine scenarios in which replacement increases good for beings and does not lead to any additional suffering for existing beings.

While there might be something to these possible explanations, recognizing them probably does not "cure" the anti-replaceability intuition for most people. And perhaps the anti-replaceability intuition should not be cured. In the next section, I will consider proposals from Cohen and Dworkin that oppose replaceability on axiological grounds.

3.2 Cohen's and Dworkin's Axiological Objections to Replaceability

ERS entails that value creation, value conservation and value destruction are not indifferent matters, but that the choice between value conservation and value-destruction-plus-creation is axiologically indifferent if the replacement value is equal in magnitude to the value it is replacing. Dworkin and Cohen both claim, in contrast, that for certain kinds of value, value creation can be indifferent while conservation of existing bearers of that kind of value is not. In (1993), Dworkin argues that the intrinsic value of human life is worth preserving once it exists, but that this is not the sort of value we should create just so there is more of it. In (2013), Cohen applies similar reasoning to institutions and objects, hinting that the reasoning might be relevant for living beings as well.[14]

Cohen's and Dworkin's proposals do not seem well suited as potential modifiers to ERS, since they are centered on non-personal value, and ERS is concerned with personal value. But their arguments are worth considering as axiological attempts to

[14] Cohen seems to be more interested in how his argument applies to the non-sentient realm because he thinks that deontological considerations protect existing human value.

avoid replaceability, and they bear some resemblance to the deontological complement to ERS that I will propose in 3.5.

Cohen defends what he calls "little c- conservatism." He writes, "Conservative conviction, as I understand it…exhibits a bias in favor of retaining what is of value, even in the face of replacing it by something of greater value no matter how much greater its value would be" (149). And, "Conservation of what has value is not, in the conservation of energy sense, conservation of value, for you lose no value, value *itself* is conserved, when you destroy something valuable and replace it by a thing of the same value. The conservative disposition is not to keep the value rating high but to keep the things that now contribute to that rating" (153).

I see two main problems for Cohen's view, aside from its ERS-incompatible acceptance of non-personal moral value. One is that without further argument, it implies an implausible view about disvalue – that we have more reason to destroy already existing disvalue than to prevent new and greater disvalue from arising. Cohen writes, "Wanting to conserve what has value is consistent with wanting to destroy disvalue" (172). If already existing positive value is prioritized for conservation and continuation over merely potential positive value that would be greater in magnitude, then unless Cohen can show that there is an asymmetry between value and disvalue in this regard, it looks as if existing disvalue is prioritized for destruction over potential disvalue that would be greater in magnitude. Yet most of us think it is more important to reduce greater badness by preventing the creation of a horrible life than to reduce lesser badness by helping an existing life be slightly less bad. To use an example more

consistent with those Cohen discusses, Cohen's view (absent a value/disvalue asymmetry) entails that it would be more important to destroy existing injustice than to prevent a new injustice from arising, even if the new injustice would be greater in magnitude.

The other problem with Cohen's view is that old value bearers are constantly destroyed and replaced with new value bearers, and if there is an irreplaceable loss in the destruction of the original value bearers that is not compensated for by new value bearers of greater value, this means things are constantly getting worse even if value is increasing. Cohen accepts this implication. He writes, "Anticonservatives say, 'But people have always said that things are getting worse,' and anticonservatives mean thereby to convey that the conservative lamentation expresses an illusion. But it is entirely possible that at any rate certain *kinds* of things have *always* been worse than they were before" (144). I call this Implausibly Accumulating Worseness and discuss it in the next section. If Cohen's view is symmetric with regard to disvalue, it also implies Implausibly Accumulating Betterness – that things could be getting better over time even if disvalue is becoming greater overall. In a straightforward interpretation of Cohen's view, this could happen if existing disvalue bearers are destroyed and then replaced by value bearers of greater disvalue. With Implausibly Accumulating Betterness, we can start with one slightly bad life which is replaced by slightly worse lives until we get to a tortured life of abject misery, and we have to say the state of affairs involving this tortured life is better than the state of affairs would have been had the original slightly bad life persisted. This results from the claim that it is better to

destroy existing disvalue than to prevent the emergence of new and greater disvalue. These results are highly unintuitive.

Dworkin's reasoning is similar to Cohen's, but Dworkin is particularly interested in the intrinsic value or sanctity of already existing human lives. Dworkin writes, "Recognizing the sanctity of life does not mean attempting to engineer fate so that the best possible lives are lived overall; it means, rather, not frustrating investments in life that have already been made. For that reason, liberal opinion cares more about the lives that people are now leading, lives in earnest, than about the possibility of other lives to come" (99). He sees the politically conservative opinion as similar, except that it is often more concerned with protecting the biological process that begins with the joining of a sperm and egg, which might have religious connotations as being related to God's creative powers. Dworkin believes that the creation and growth of a human life is something that should be maintained once it is started; the process never starting is not bad, but the process starting and then prematurely ending often is. Dworkin writes, "No frustration of life is involved when fewer rather than more human beings are born, because there is no creative investment in lives that never exist. But once a human life starts, a process has begun, and interrupting that process frustrates an adventure already under way. So the idea that we deplore the frustration of life, not its mere absence, seems adequately to fit our general convictions about life, death, and tragedy" (88).

Dworkin's point here may help to explain rather than justify intuitions that something worse happens when someone dies and is replaced by a happier individual. The intuition that there is an important axiological distinction between creation of a new

life and extension of an existing life could in part be due to the waste and trauma involved in destroying an existing being in order to create a new one, and the greater work involved in creation as opposed to in extension. We might have been less inclined to think of creating and extending as so different if everyone were born as fully formed, mentally capable, independent adults. We might be less inclined to think it an obligation to save children drowning in ponds if we then had to raise all of those children into adulthood. The Reconciled Epicurean agrees with Dworkin that a process is set into motion at conception and that the biological and creative investment into a life is wasted with a premature death, and that this counts against replacement. But if Dworkin thinks the creation of a new life – no matter how valuable – cannot make up for the loss of the previous life, then his argument, like Cohen's, also runs into the problem of Implausibly Accumulating Worseness. I will discuss this problem next. I will also discuss a problem with anti-replaceability that I call Axiologically Irreplaceable Mediocrity.

3.3 Accumulating Worseness and Axiologically Irreplaceable Mediocrity

Imagine that in her wanderings around the universe, God comes across a planet in which there is one being with a life that is just slightly more bad than good. God dubs this the Barely Bad Life. Given her limited powers, God has three options. If she does nothing, this life will persist in its just barely bad state indefinitely. Another option is for God to replace Barely Bad Life with Barely Bad Life 2, which will be a new life of the same quality and duration as the original Barely Bad Life. God's third option is to replace Barely Bad Life with a being who will also live indefinitely, but have just

slightly less badness in its life than Barely Bad Life does, so that it would be axiologically indifferent whether this life continues, as opposed to being slightly bad for the life to continue. This would be Neutral Life. None of these potential alternative lives would have any cultural or genetic connection with the first Barely Bad Life, and so there is no benefit of meaningfulness that Barely Bad Life could receive from the creation of anyone else.

The first two options of either leaving Barely Bad Life to persist, or replacing it with Barely Bad Life 2, appear to be equally bad, though there is a case to be made that replacing Barely Bad Life with Barely Bad Life 2 is an improvement in that it spreads out badness amongst more individuals. This possibility raises complex questions about the aggregation of value across individuals that I do not have the space to explore here, so I will assume that the replacement of Barely Bad Life with Barely Bad Life 2 is approximately equal in disvalue to letting Barely Bad Life persist. Leaving Barely Bad Life to persist is just slightly overall bad for that life, and ending it would end this badness. However, if this is done through replacement, the disappearance of badness for Barely Bad Life is outweighed by this addition of the same level of badness for Barely Bad Life 2. In contrast, replacing Barely Bad Life with Neutral Life would end badness for Barely Bad Life and put in its place someone whose life has a little less badness. God therefore decides to replace Barely Bad Life with Neutral Life.

Now the one life on this planet has a particular balance of good and bad that makes it axiologically indifferent whether this life continues. God has three new options. God can do nothing and allow Neutral Life to persist neutrally indefinitely.

Another option is for God to replace Neutral Life with Neutral Life 2, which will be a new life of the same quality and duration as the original Neutral Life. God's third option is to replace Neutral Life with Barely Good Life, which will be a life of the same duration as Neutral Life, with the difference being that Barely Good Life has a little more good in its life than does Neutral Life – enough so that Barely Good Life is just barely worth living as opposed to its continuation being a matter of indifference.

The first two options seem to be equally neutral. There is no objectionable value disappearance if God replaces Neutral Life with Neutral Life 2, as the balance of good and bad for Neutral Life is axiologically indifferent, as would be the balance of good and bad for Neutral Life 2. In contrast, the potential Barely Good Life will be worth living, albeit just barely so. It seems there will be more good for beings if God replaces Neutral Life with Barely Good Life, so that is what she does.

Now the planet has a life that is barely worth living. God has three new options. She can leave Barely Good Life persisting indefinitely in its life that is barely worth living. She can replace Barely Good Life with Barely Good Life 2, which a life that is no more good or bad than the original Barely Good Life. Finally, she can replace Barely Good Life with Slightly More Good Life, which will have the same duration as Barely Good Life, but will have just a tiny bit more good in it than Barely Good Life does.

Since replacement was axiologically preferred in all the previous cases in which the next life would have slightly more good than the previous life, we might now suppose that God ought to replace Barely Good Life with Slightly More Good Life. But even though Barely Good Life is just slightly better off than Neutral Life, this slight

difference may be significant. Since Barely Good Life is worth living, its destruction is objectionable. This was not obviously the case for Neutral Life or Barely Bad Life, at least if our considerations are purely axiological. Perhaps the value in creating Slightly More Good Life does not make up for the destruction of Barely Good Life.

Yet there is something puzzling about this possibility. We have no problem measuring the value of a state of affairs by just looking at the value of one life and comparing that to the value of the replacing life, so long as none of the replaced lives are worth living. We see the destruction of one bad or neutral life for another bad or neutral life that has the same mixture of bads and goods as preserving the same level of value. Exchanging one bad life for another bad life of the same magnitude is equally bad as not making this change. Exchanging a neutral life for another neutral life is as neutral as not making the change. We do similar math when we replace bad or neutral lives with lives that have slightly more good than the previous life. When we destroy Barely Bad Life for Neutral Life, we judge this as a slight improvement that is equivalent to the difference in goodness between Barely Bad Life and Neutral Life. When we destroy Neutral Life for Barely Good Life, we judge this as a slight improvement that is equivalent to the difference in goodness between Neutral Life and Barely Good Life. But as soon as we are thinking of exchanging Barely Good Life for Slightly More Good Life, this math possibly stops working. Some opponents of ERS who see no problem with exchanging bad or neutral lives for equally bad or neutral lives, or exchanging these for more good lives, cannot say that exchanging Barely Good Life for Barely Good Life 2 is equally good as not making the change, and cannot say

that exchanging Barely Good Life for Slightly More Good Life improves the state of affairs for beings by the difference in goodness between Barely Good Life and Slightly More Good Life. Why does the math stop working?

One answer the ERS opponent can give is that no amount of goodness in a replacing life can make up for the destruction of goodness for the replaced life. This answer leads to Axiologically Irreplaceable Mediocrity (as well as Implausibly Accumulating Worseness). With this answer, even replacing Barely Good Life with Delightful Life – a being with an indefinitely persisting life of bliss – is less valuable than letting Barely Good Life persist.

If God is discouraged from replacing Barely Good Life with Delightful Life on axiological grounds, then it must be that a tiny amount of good for an existing individual is worth more than any amount of good there would be for someone who could exist but does not yet exist. McMahan might claim that the difference is that the continuation of Barely Good Life provides a comparative benefit to the individual with that life, whereas the creation of Delightful Life would bestow a non-comparative benefit on the individual with that life, and that a comparative benefit is more important than a non-comparative one. But I have argued that extending a life provides a non-comparative benefit rather than a comparative one, and that anyway we should not favor comparative value over non-comparative value of the same magnitude. Such a steep value discount for newly created lives – or perhaps steep value premium for already existing lives – does not seem consistent with the notion that there can be value in creation of lives and not just in extension of lives. Unless we want to deny that there is

value in creating lives – and I have already tried to show that we should not deny the value in creating lives – it seems we must say it is more valuable for God to replace Barely Good Life with Delightful Life.

Another implication of answering that no amount of goodness in a replacing life can make up for the worseness of death for the replaced life is Implausibly Accumulating Worseness. When generations of beings have lives that are worth living, but which eventually end, a world's state of affairs gets worse over time as generations replace each other, even if the lives of each generation are as good as the lives of the previous generation, or are better than those of the previous generation. Imagine life on a planet starts with a Barely Good Generation and becomes slightly more good for each proceeding generation until we arrive at a generation that has lives of pure bliss. It looks as if we may have to say that the state of affairs is worse once we arrive at the bliss generation than it was in the early history of the planet when there was the first generation of beings with lives that were just barely worth living, and that we should have wished for that first generation of barely worth living lives to persist indefinitely instead of being replaced by steadily improving lives. McMahan makes a point similar to this when discussing the claim that there is a loss of impersonal value when a species goes extinct which cannot be compensated for when a new species comes into being, and notes the troubling link between the ever-accumulating worseness of species extinction and species replacement with the ever-accumulating worseness of individual death and replacement:

Since animals first appeared, an indefinite number of species have become extinct and an indefinite number of new species have arisen. If the extinction of a species involves a loss of impersonal value that cannot be made up for by the appearance of a new species, it seems that the world must have got worse with every extinction, even when the extinction of one species has coincided with the appearance of more than one new species. It thus seems that, according to this understanding of the value of species, the world would have been better impersonally if none of the earliest species had become extinct, even if a consequence of that would have been that most of the newer species, perhaps including the human species, would never have existed.

This conception of the impersonal value of species also raises the question how the extinction of a species could involve a loss of impersonal value when the coming-into-existence of that species did not involve a gain in impersonal value. A parallel and perhaps more familiar question arises if we think…that individual human lives are sacred. For while we seem to think that the death of a person involves the loss of something intrinsically valuable, most people do not think that it is impersonally better, or that there is a gain in value, whenever a person comes into existence. Although these do seem…to be commonsense beliefs, they do seem to be in tension with one another. (2016:276-277)

Implausibly Accumulating Worseness implies that God may replace Barely Bad

Life as many times as she wants, so long as she exchanges it with other barely bad lives, and she may replace Neutral Life as many times as she wants, so long as she exchanges it with other neutral lives, but she cannot exchange lives worth living with lives that are even more worth living. She can make 1,000 life replacements so long as the original life is not worth living, and the lives of these replacements never increase in value in comparison to the previous life, but she cannot replace Barely Good Life with increasingly happier lives until she arrives at Delightful Life.

I believe this shows we should reject the claim that no amount of positive value in a replacement life can make up for the destruction of good for the previous life. However, this does not show that God is axiologically justified in replacing Barely Good Life with Slightly More Good Life. It might be that the replacing life needs to be far more worth living than the replaced life to make up for the destruction of good for the previous life.

Suppose God has decided to replace Barely Good Life with Delightful Life, but that there are two ways she can do this. With the first method, there will be a 10,000-year gap between the end of Barely Good Life and the beginning of Delightful Life. With the second method, she can do the replacement with 1,000 decade-long lives in between, each of which will be slightly more good than the last. The first of the 1,000 lives will have a little more good than Barely Good Life, the next will have more good than that, and so on, until we reach Delightful Life. All the decade-long lives in these 1,000 generations will be worth living, since even the Barely Good Life is worth living, but there will also be a destruction of their good at their deaths. If this destruction of

good is enough to outweigh the benefit of life for the replacing life that is just slightly happier, then it looks as if we cannot endorse the extra thousand generations in between Barely Good Life and Delightful Life, and we should prefer that God choose the replacement option with the lifeless 10,000-year gap instead.[15]

It seems to me that it would be far less valuable to have a 10,000-year lifeless gap between Barely Good Life and Delightful Life than to have a thousand decade-long gradually improving lives on the way to Delightful Life. However, if we do think the most valuable option is for God to cycle through 1,000 one-decade lives that slightly increase in value on the way from Barely Good Life to Delightful Life, then we seem to accept the axiological replaceability of beings even when the replacements have just slightly more value than those being replaced.

Those who think this goes too far and want to deny ERS in order to say that the destruction of a value bearer and a value bearer's good at death is more worthy of avoidance than the good of life is worthy of creation seemingly have to accept that death can be a reason not to create life, and they have to accept Implausibly Accumulating Worseness and Axiologically Irreplaceable Mediocrity. However, if they think that the disappearance of a value bearer and a value bearer's good at death is not always more worthy of avoidance than the good of life is worthy of creation, because they think God should replace Barely Good Life with Delightful Life – but they also think the replacing life cannot be just slightly more valuable than the life being replaced to axiologically justify the replacement – then they face the challenge of identifying and

[15] This example relates to Temkin's spectrum arguments in *Rethinking the Good,* but I do not have the space to draw out the parallels and consider the relevance of his arguments here.

defending the amount of value difference there must be between the replaced life and the replacing life to justify replacement. They also must accept Implausibly Accumulating Worseness in most realistic scenarios, because disvalue would steadily increase in a world in which each generation of beings have lives just slightly better than those of the previous generation. For a world to be worthy of creation, each generation would need to have lives that are much better than those of the previous generation, but if there is any upper limit on the feasible goodness of a life, such a world would start accumulating worseness before too long. A world of beings who are at near maximum happiness who breed and die and are replaced by their offspring who are also at near maximum happiness would just get worse and worse as time goes on, even though the quality of the lives would be staying exactly the same. I believe ERS offers a more plausible assessment than this.

However, something may still seem unsettling about the replaceability that ERS cannot prohibit on its own. It is therefore worth considering a modification to ERS to avoid replaceability in practice, even if it cannot be avoided in our axiology. In the next section, I will consider and reject a non-axiological barrier to replaceability that Velleman and his colleagues have proposed.

3.4 Velleman et. al.'s Non-Axiological Objection to Replaceability

Velleman, Darwall and Anderson make a claim that implies a non-axiological objection to replaceability. The claim is that value matters only derivatively, and only if the bearers of the value matter. While the view is illuminating, I will argue that it is not an

appropriate ERS-friendly buffer against replaceability, though it can help to inspire one.

Anderson writes, “Persons, not their interests, bear intrinsic value” (1993:225). Velleman writes, “[W]hat’s good for a person is worth caring about only out of concern for the person, and hence only insofar as he is worth caring about. A person’s good has only hypothetical or conditional value, which depends on the value of the person himself. … [O]ur appreciation of value that accrues to someone depends on a prior appreciation of a value inhering in him” (2015:27). This is “interest-independent value,” and Velleman says that it relates to Kant’s concept of “dignity.”

Frick has put this idea to use in population ethics; he believes it helps him defend the conjunct of the population asymmetry which holds that there is no moral reason to create new lives, however good these would be. He writes, “Human well-being matters because people matter – *not vice versa*” (351). Since non-existent possible people are not persons or value bearers, their possible future wellbeing does not matter if we do not create them, and so we need not create them just so they can have wellbeing. Velleman and Anderson also see this idea as having similar implications for population ethics. Anderson criticizes the consequentialist view that cares about value rather than value bearers, saying this means that respecting humanity, “becomes, absurdly, a matter of increasing the number of human beings, as the population paradoxes show” (71). Velleman writes, “Dignity is what Kant called a ‘self-existent’ value — a value to which we are obliged to respond only when it already exists, and then only by paying it reverence or respect. The value of persons does not oblige us to maximize the number of people in existence; it obliges us only to respect the people

who do exist" (33).

This view would prohibit replaceability because it would offer protections for the currently existent without placing any value on the creation of new lives. A problem with this view, and possibly even with Frick's population-ethics-centered handling of it, is that if wellbeing only matters if people matter – and if we interpret this to mean that the possible future wellbeing of a non-existent potential individual does not matter – it seems we have no reason not to create miserable lives. Velleman is less concerned with misery than a lack of dignity in lives, so we might instead say that if we accept the Anderson/Darwall/Velleman/Frick view, there is no reason not to create undignified lives for the sake of the people who will therefore be subjected to undignified existences.

Velleman's elaboration of the view is problematic for the currently existent as well. Since people matter categorically and their wellbeing matters only conditionally, Velleman argues it is immoral for a person to kill herself, even if her life has become intolerable – so long as she still has her dignity, and thus matters as a person. Velleman thinks suicide can be appropriate once living with dignity has become impossible for them. But consider someone who is in extreme pain for much of her life, but manages to live with dignity. We can imagine that she is in pain for most of the day, but for about two hours of each day, she is able to exercise her rational faculties and work on projects. We can also stipulate that when she is in pain, she is not screaming and writing or embarrassing herself in any way. Instead, she sits in a chair, still, with her eyes closed, quietly suffering the most extreme pain anyone has ever felt. Velleman's view

entails that it would be immoral for her to end her life if she no longer wanted to endure a life of torture. Though Velleman does not quite put it this way, the claim seems to be that those who have lost their dignity have lost what makes them matter as people – they have lost their interest-independent value – and thus have lost what grounded the prohibition against ending their own lives. It is difficult to see how else the suicide prohibition could be lifted in Velleman's view, but if someone's ceasing to matter as a person is what allows them to end their own lives, the view that value only matters if the value bearers matter suggests that it is permissible for anyone to torture someone who has lost their dignity. Perhaps Velleman would respond that someone who once mattered as a person always retains some vestige of the interest-independent value that prohibits such torture, but then it seems he must prohibit their suicide as well. Furthermore, we can imagine sentient beings who never have interest-independent value – perhaps because they are not rational, or they live purely in the moment – and on this view their suffering would not matter because they do not matter. Yet it seems that if we can costlessly stop outrageous pain that is instantiated in a non-mattering being, we ought to stop it.

I believe these considerations refute the Anderson/Darwall/Velleman/Frick view.

3.5. Avoiding Replaceability Without Abandoning ERS

ERS is an axiological theory. If it is right, it gives us a general way to think about the value of conscious life – its creation, its non-creation, its continuation, its destruction. If ethics is nothing more than a direct application of axiology, then ERS entails

replaceability. But there may be more to morality than a value-based evaluation of outcomes. This is suggested by the common intuition that it is wrong to murder an innocent individual to save five individuals, even if we do not think that killing one to save five produces less good than letting the five die, and even if we do not think it is wrong for someone to choose to sacrifice themselves to save five.

Willingly sacrificing one's own life to extend other lives is a form of replacement. A unique and irreplaceable individual is lost with this sacrifice, but many find it admirable and sometimes even obligatory for one individual to give their life up to save one or more individuals in certain circumstances. Here we have no problem accepting that a unique individual can be traded for one or more other unique individuals. But there are two important differences here from the sort of replaceability that I have said the bare ERS entails. The first is that no one has forced this individual to make the sacrifice. The second is that the sacrifice is for beings who already exist. Perhaps surprisingly, the proponent of ERS can accept that these differences are morally important and thereby avoid replaceability. Since it is the second difference that may seem most at odds with ERS, that is the one I will address for the rest of this chapter.

ERS entails that there is no value difference between the possible future good of an already existing being and the possible future good of a potential being. But it need not entail that there is no moral difference between assuring the good of one over the other. Axiologically, ERS looks like the impersonalist total view in a wide-person-affecting guise. But the person-affecting nature of ERS, combined with its denial of the personal comparability of existence and nonexistence, allows it to draw a deontological

distinction that the impersonal value maximizing totalists seemingly cannot. The nonexistence of merely potential beings allows us to deny them moral standing. We could say that we do not have obligations to possible beings to create them, but we do have obligations to created beings, and to beings who will be created. These obligations can include protecting their lives if these are good lives. When these protections of the existent conflict with the creation of a new being, the conflict is resolved in favor of those we have an obligation to: the existent. And this is true even if the existing life is just slightly good and the life this slightly good life prevents us from creating is a delightful life. At the same time, we can acknowledge that creating the new life would be very good, and would in fact be more good than allowing the existing life to continue instead.

Let us say that we have no obligations to the merely possible, we cannot wrong the merely possible, and we cannot do right by the merely possible. In contrast, we do have obligations to the existent, we can wrong the existent, and we can do right by the existent. While someone's death is experientially and axiologically the same as their non-creation, it need not be morally the same as their non-creation, because killing or allowing the death of an existing individual could wrong them, while non-creation cannot wrong anyone.

Once we create someone, they are an actual being. It is therefore possible to right or wrong them through their creation. If coming into existence is good for someone, we can say that by creating a good life, we do right by the being who comes into existence and benefits. Yet it would not have been wrong to not create this being,

since in that case there would have been no one we could wrong. If coming into existence is bad for someone, then creating someone with a bad life wrongs this being, even though not creating them would not have been right since then there would have been no one to do right by. In short, creating a good life is right and creating a bad life is wrong, but not creating a life is neither right nor wrong because there is no one whom can be righted or wronged by non-creation.

It may now look as if I am about to face a familiar population ethics dilemma. We might have thought the rightness of creating a good life gives us an obligation to create that life, and the wrongness of creating a bad life gives us an obligation not to create that life, in which case we end up with something like the impersonalist total view, and it is again unclear how we are to avoid replaceability. Alternatively, we might think the rightness of creating a good life is not the sort of rightness that entails an obligation. But then the wrongness of creating a bad life must not be the sort of wrongness that entails an obligation either, and so we lack an obligation to avoid creating bad lives. We therefore end up having to embrace and rationalize the problems with typical narrow person-affecting views instead.

I believe we can avoid this dilemma. Applying ERS claims to creation and non-creation suggests a perhaps unusual deontological category of not-necessarily supererogatory cases in which we can do something right when not doing this right thing would not have been wrong, and cases in which we can do something wrong when not doing this wrong thing would not have been right. We can utilize this category to avoid replaceability while maintaining a prohibition on creating bad lives with the

following asymmetry that even the symmetry-partial Reconciled Epicurean can accept:

The deontological asymmetry: We have an obligation not to do what is wrong, even if avoiding the wrong thing would not be right, but we do not have an obligation to do what is right if not doing the right thing would not be wrong.

If this asymmetry is plausible, it seems to address the previous concerns. We have an obligation not to create miserable lives because doing so would wrong the individuals born into the miserable existences, even though not creating the miserable beings would not be right, since we cannot do right by the never existent. It would be right for us to create happy lives when doing so does not wrong anyone, but since it would not be wrong to not create these happy lives, we are under no obligation to create them. Since it would be wrong to destroy existing good lives, even if the lives are just slightly good, and not wrong to not create blissful lives – despite the rightness of creating blissful lives – it is impermissible to destroy a slightly good existing life in favor of a possible blissful life.

Accepting ERS, deontological considerations in favor of the existent over the non-existent, and this deontological asymmetry allows there to be good and bad value in creation, a prohibition on creating bad lives, as well as reasons in favor of creating good lives that do not rise to the level of an obligation — with these reasons increasing in strength the more good the new life would be. Doing what is right and doing what is good is not a matter of indifference even when these are non-obligatory. To the extent

that we are motivated by these moral reasons, we will create the best lives possible if we create lives at all. It would not be violating a moral obligation to create a life that is barely worth living when we can create a life that is much more worth living, but it would flout the stronger moral reasons we have to create the life that is much more worth living. It would be more right to create the more good life, so if we chose to create the mediocre life instead, we would presumably need prudential reasons or conflicting moral reasons weighing in favor of the mediocre life to justify the choice. Furthermore, the obligation we have to existing lives means that we may often have obligations to help existing individuals have better lives, and these obligations will be easier to meet if we create individuals in circumstances that we expect to be very worth living in as opposed to barely worth living in.

This may seem to categorize the creation of good lives as supererogatory. However, supererogation usually relates to demandingness, and demandingness is not why creating blissful lives is optional in the view I just presented. Creating blissful lives is optional because there is no one to whom creation can be an obligation, and we cannot right or wrong the non-existent. This might suggest a need for another category of supererogatory acts that are praiseworthy but not required, without necessarily being demanding. Or perhaps a completely different term would be more appropriate.[16] Regardless, if there can be moral reasons that do not rise to the level of an obligation because the act is too demanding, it seems there can also be moral reasons that do not rise to the level of an obligation because the potential beneficiary does not exist. And so

[16] It is possible that such a term has been proposed, and that I am not aware of it.

we have a prohibition against creating bad lives, as well as moral reasons in favor of creating good lives that provide practical moral direction but not obligation.

This view I have proposed has straightforward implications for replaceability and our obligations to future beings.

To consider the implications for replaceability, imagine there is a machine that destroys the being who walks into it, and creates a replacement being with many of the same memories and personality traits as the original, but who is much happier, more effective and more helpful than the original. The replacement being who steps out of the machine takes over the life of the original without any problems, but does so while being happier, nicer to be around, and contributing more to the world than the original would have. ERS combined with the deontological considerations I have proposed entail that it is permissible to step into this machine (because there is only personal value, so this does not destroy interest-independent value), it would not be bad or worse for the individual who steps into the machine (because death is not bad or worse for the person who dies), it would be doing something right and more good to step into this machine (because it would create a happier and more helpful being), but it is not obligatory to step into this machine (because there is no obligation to right someone by bringing them into existence), and it would be wrong to push someone into the machine (because we have obligations to the existent).

I will now consider the implications this view has for human extinction and our obligations to future beings. If we all decided to stop having children, we would be failing to do what is more good and failing to do what is right for the individuals we

could have conceived, but we would not be violating any obligations and we would not be doing anything wrong. However, we can have obligations to beings in the future who will exist, either because we had some causal influence in their creation or because they will exist independently of our actions. Assuming we have some causal influence in their existences, then we do what is good for them and what is right for them by helping them to exist in good lives, but we can still wrong them by depriving them of resources because of actions that damage the environment, for instance. If our actions now lead to lives in the future being barely worth living, we will have contributed to the goodness and rightness of the coming into existence of these specific lives, but we will have done less good and less right than if we had helped assure that lives of the future would be very much worth living.

I am aware of four potential problems for the deontological addition to ERS that I have proposed. I will address these next.

Problem 1. This view leads to deontologically irreplaceable mediocrity. I argued that it was a problem for Cohen and Dworkin that their view led to Axiologically Irreplaceable Mediocrity, but the deontological apparatus I have attached to ERS means that if God comes across a planet that is hosting a life that is barely worth living, it would be wrong for her to destroy this life to create the delightful life, just because the barely worth living life got here first. If this is a real problem, it is easily resolvable: we can abandon my deontological suggestion, translate the bare ERS directly into reasons for action and accept replaceability. My deontological suggestion is for those who think replaceability is wrong, but not everyone agrees that it is, and they will want to reject

my deontological suggestion. Even those who advocate deontological replaceability should like that at least the modified ERS does not entail Accumulating Worseness as lives die and are replaced by better lives. It would be wrong for God to destroy Barely Good Life for Slightly More Good Life, but it would not be bad or worse for beings if she did this.

Problem 2. This view entails that we may not create wonderfully blissful lives if we must commit a tiny wrong in order to do so. It is wrong to flick someone on the ear without their consent, but it is not wrong to fail to create a billion happy lives on Planet Zog. Imagine the only way we can create a billion happy lives on Zog is by flicking someone on the ear without their consent. My deontological replaceability prevention strategy seems to suggest that we should not flick the ear despite all the good it will cause.

This may be a bullet I must bite, or it might be possible to fine-tune how we weigh the obligations we have not to commit minor wrongs against the moral reasons we have in favor of bestowing substantive goods. Perhaps it is permissible to flick someone on the ear to create a billion happy lives because the magnitude of the goodness and non-obligatory rightness of creating all these happy lives so vastly outweighs the magnitude of the minor wrong of the ear flick, but it still remains impermissible to kill a happy innocent individual to create all the blissful Zog lives. Working out the strength and relevance of the non-obligatory moral reasons in favor of creating happy lives is tricky, and I do not have the space to consider this further here.

Problem 3. On this view, it is morally safer to not create lives at all when there is

some risk of the lives being bad. Imagine that if we press a button, there is a 95 percent probability that a billion very happy lives and no bad lives will materialize on Planet Zog, and a 5 percent probability that one bad life and no happy lives will materialize on Planet Zog. It is very good and right to press the button if it will create the happy lives, and it is wrong to press the button if the bad life appears instead. We are not obligated to press the button if it will create the happy lives, but we are obligated to not press the button if it will create the bad life. My view seems to entail that we should not press the button. The total view, in contrast, has us press the button because of the expected value of doing so.

This problem is similar to the previous one, except that it involves the risk of wrongdoing instead of a guarantee of wrongdoing. This seems to make a difference. While my proposal does not permit a straightforward expected value calculation for all our moral decisions, since this would lead to replaceability, it does seem that we can consider the weight of the wrong we risk committing against the weight of the good and the right that we are hoping to achieve. We take risks of committing a wrong all the time, and it is arguably not impermissible to do this, especially if we are aiming to achieve a greater right. When we drive a car, we take a risk of wrongly running someone over, yet we often drive cars anyway, even for frivolous reasons. While we are not obligated to press this button that will hopefully create the billion happy lives but possibly create the one bad life on Zog, and we might be obligated not to press it, there are very strong moral reasons in favor of pressing the button, and the right and good we are hoping to achieve by pressing it are much stronger than the wrong we risk by

pressing it. I believe that under my proposal, it is permissible and indeed highly encouraged to press this button.

Problem 4. ERS is incompatible with the claim that killing is wrong in a person-affecting sense because for there to be a wrong, someone must be wronged, but ERS entails that there can be no one who is wronged by death. ERS therefore cannot escape replaceability no matter what deontological apparatus we attach to it. An Epicurean cat sitter is hired by a cat owner to take care of her cat named Whiskers. She tells the cat sitter that he has a duty to Whiskers that he keep Whiskers alive. About halfway into his assignment, the cat sitter puts some painlessly fatal cat poison powder in Whiskers' food. Whiskers eats the powder and painlessly dies. When Whiskers' owner returns and sees that Whiskers has died, she accuses the cat sitter of having failed his duty to Whiskers that he keep Whiskers alive. The Epicurean cat sitter retorts, "When Whiskers was alive, I had not failed in my duty to keep him alive. When Whiskers died, there was no Whiskers, and I cannot have a duty to the non-existent." If this claim is consistent with ERS plus my deontological addition to it, ERS seemingly cannot accommodate a moral prohibition against killing. It could still be the case that we cannot seriously wrong someone in order to create a new being. For instance, it would be impermissible to torture someone to create a new being. However, it would not be impermissible to kill someone to create a new being, because this would not wrong the being who was killed, which would mean there is no block against replaceability after all.

There may be a solution. While we cannot have a duty to the never existent, we can have duties to the temporarily existent that arise during their existence. When

Whiskers was alive, the cat sitter had a duty to Whiskers that he keep Whiskers alive. In my proposal, the cat sitter admittedly cannot have any duties to Whiskers when Whiskers no longer exists, but this does not change that the cat sitter did have a duty to Whiskers to keep him alive. It is plausible that he fails in this duty the moment Whiskers dies. The duty to keep Whiskers alive that existed when Whiskers was alive is not retroactively dissolved because Whiskers is no more.

If my responses to these objections have been adequate, and I have not overlooked any others, then it seems we can have the benefits of the ERS axiology without permitting the murder of existing people in order to create mildly happier people.

Conclusion

In this thesis, I have argued for a view that was conceived of by a philosopher who did not himself endorse it, and which has not been discussed much sense. This is the Epicurean Reconciliation Strategy (ERS). McMahan devised it as a compromise position for those convinced by the Epicurean argument against the badness of death, but who were disturbed by its seemingly radical implications for ethics and value theory. Epicurus' argument is simple and compelling: death is not bad for the one who dies, because if death is the obliteration of the subject, there essentially is no one who dies. But this is often interpreted as leading to full neutrality about existence – indifference as to whether lives are created and whether lives persist, no matter how good or bad the lives are or would be. However, McMahan showed there is no contradiction in claiming that life can be good and worth creating and continuing even if Epicurus is right that nonexistence is not bad. McMahan's brief sketch of the view left plenty of room for elaboration and defense, and I have attempted to fill the gaps.

I began by developing a conceptualization of nonexistence to support the claim that existence and nonexistence are not comparable in personal terms. We cannot compare existence and nonexistence this way because nonexistence simply offers nothing for a comparison. This means we should reject views in population ethics which hold that it can be better or worse for someone that they have come into existence. Most population ethicists are happy to reject such views. But then I went further and claimed that this personal incomparability of existence and nonexistence also implies that death

is valueless, and that continuing to exist is personally incomparable with ceasing to exist. These are much more contentious claims. If I am right, death is neither good, bad, better nor worse for the person who dies. Furthermore, continuing to exist is neither better nor worse than ceasing to exist, just as coming into existence is neither better nor worse than never existing. This leads to a symmetric stance about the value of creating lives and the value of continuing lives, rendering implausible the position that there is no value in creating new lives, as this entails the valuelessness of extending existing lives. While full neutrality about existence allows that we should improve a life of pure suffering if we can do so within the life, it is unable to allow that suffering is a reason to end a life, because we cannot compare suffering and nonexistence. Full neutrality about existence thus implies that we should flip a coin if we are choosing between nonexistence and an eternity of torture. It is difficult to believe that anyone would endorse this. Even the most enthusiastic advocate of neutrality about existence would choose nonexistence over eternal suffering "just in case" they might be wrong about the impossibility of making a rationally justifiable choice between the two.

Fortunately, we can escape this full neutrality about existence by recognizing the non-comparative, intrinsic, personal value of good and bad that is not comparatively better or worse for anyone, and by acknowledging reasons to create, extend and destroy lives on the basis of this non-comparative value. I showed that this non-comparative value exists, is action guiding, and is never less important than comparative value.

These arguments led us to ERS – a view with many appealing features, but with a major public relations problem: the replaceability of rational beings appears to be an

obvious implication of ERS. I argued that from a purely axiological perspective, there is no plausible escape from replaceability. But there is a plausible deontological escape from it. I considered and rejected Cohen's, Dworkins', and Velleman et. al.'s answers to replaceability before proposing my own deontological apparatus which is consistent with ERS. This modified ERS has many advantages. It allows that there is value in creating lives and in extending lives. It entails that it is bad to be born into a bad life, and good to be born into a good life. Because there is value in coming into existence, and no negative value in going out of existence, the world is not necessarily getting worse as generations replace each other. Instead, the value of the world depends on the value of the lives in it. However, replaceability is impermissible because we have obligations to the existent and not to the non-existent. The deontological asymmetry I proposed allows us to say that since it would be wrong to create bad lives, we are prohibited from creating bad lives despite lacking obligations to the non-existent, but we are not required to create good lives even though it would be right to do so, because we are not required to do what is right if not doing it would not be wrong. Even though creating good lives is not required, it is not neutral and indifferent either: it is good, right and praiseworthy.

There is much left to consider. ERS has implications for issues in applied ethics, such as animal farming and abortion. The aggregation of value across lives and within lives is extremely relevant to how ERS is understood and implemented, and can help us answer whether ERS is committed to the Repugnant Conclusion. Other interesting questions in this vicinity include whether there can be post-mortem harms, and how

different personal identity theories and value theories such as hedonism or objective list theories interact with ERS. I hope to explore these topics in the future.

WORKS CITED

Anderson, E. 1993. *Value in Ethics and Economics*. Cambridge, Massachusetts: Harvard University Press.

Arrhenius, G. & Rabinowicz, W. 2015. "The Value of Existence." In I. Hirose & J. Olson, eds. *The Oxford Handbook of Value Theory*, 424-443. Oxford: Oxford University Press.

Bader, R. Unpublished Manuscript A. "Neutrality and Conditional Goodness."

———. Unpublished Manuscript B. "The Asymmetry."

Bradley, B. 2009. *Well-Being and Death*. Oxford: Oxford University Press.

Broome, J. 1999. *Ethics Out of Economics*. Cambridge: Cambridge University Press.

———. 2004. *Weighing Lives*. Oxford: Oxford University Press.

Bykvist, K. 2007. "The Benefits of Coming into Existence." *Philosophical Studies* 135: 335-362.

———. 2015. "Being and Wellbeing." In I. Hirose & A. Reisner, eds. *Weighing and Reasoning: Themes from the Work of John Broome*, 87-94. Oxford: Oxford University Press.

———. Unpublished Manuscript. "Person-Affecting and Non-Identity."

Cohen, G.A. 2013. "Rescuing Conservativism: A Defense of Existing Value." In M. Otsuka, ed. *Finding Oneself in the Other*, 143-174. Princeton: Princeton University Press.

Dancy, J. 2013. "Meta-Ethics in the Twentieth Century." In Beaney, M. ed. *The Oxford Handbook of the History of Analytic Philosophy*, 729-749. Oxford: Oxford University Press.

Darwall, S. 1997. "Self-Interest and Self-Concern." *Social Philosophy and Policy* 14: 158-178.

Dworkin, R. 1993. *Life's Dominion: An Argument about Abortion and Euthanasia*. London: Harpers Collins Publishers.

Epicurus. 1926 [around 290 BCE]. "Letter to Menoeceus." In C. Bailey, ed. *Epicurus: the extant remains*. Oxford: Clarendon Press.

Feit, N. 2002. "The Time of Death's Misfortune." *Noûs* 36: 359–383.

Feldman, F. 1991. "Some Puzzles About the Evil of Death." *The Philosophical Review* 100, 2: 205-227

Frick, J. 2017. "On the Survival of Humanity." *Canadian Journal of Philosophy* 47, 2: 344-367.

Gardner, M. 2016. "Well-Being and the Non-Identity Problem." In G. Fletcher, ed. *The Routledge Handbook of the Philosophy of Well-Being*, 429-438. London: Routledge.

Gustafsson, J. 2014. "Neither 'Good' in Terms of 'Better' nor 'Better' in Terms of 'Good.'" *Noûs* 48, 3: 466-473

Harman, E. 2009. "Harming as Causing Harm." In Roberts, M. & Wasserman, D. eds. *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*, 137–154. London: Springer.

Holtug, N. 2001. "On the Value of Coming into Existence." *The Journal of Ethics* 5, 4: 361-384.

Hutchinson, M. 2017. "People Aren't Replaceable: Why It's Better to Extend Lives than Create New Ones." In E. Gamlund & C. Solberg, eds. *Saving Lives from the Badness of Death*, Chapter 15. New York: Oxford University Press.

Johansson, J. 2010. "Being and Betterness." *Utilitas* 22: 285–302.

Kazez, J. 2011. "Absent Pains." In *In Living Color*. http://kazez.blogspot.co.uk/2011/06/absent-pains.html

Kamm, F.M. 1992. *Creation and Abortion: A Study in Moral and Legal Philosophy*. New York: Oxford University Press.

Kraut, R. 2012. *Against Absolute Goodness*. Oxford: Oxford University Press.

Lucretius. 1910 [Around 50 BCE]. *On the Nature of Things*. C. Bailey, trans. London: Oxford University Press.

Luper, S. 2009. *The Philosophy of Death*. Cambridge, UK: Cambridge University Press.

McMahan, J. 1981. "Problems of Population Theory." *Ethics* 92, 1: 96-127.

———. 1988. "Death and the Value of Life." *Ethics* 99, 1: 32-61.

———. 1998. "Wrongful Life: Paradoxes in the Morality of Causing People to Exist." In J. Coleman & C. Morris, eds. *Rational Commitment and Social Justice: Essays for Gregory Kavka,* 208-47. Cambridge: Cambridge University Press.

———. 2008. "Eating Animals the Nice Way." *Dædalus* 137, 1: 66-76.

———. 2013. "Causing People to Exist and Saving People's Lives." *The Journal of Ethics* 17, 1: 5-35.

———. 2016. "The Moral Problem of Predation." In A. Chignell, T. Cuneo & M. Halteman, eds. *Philosophy Comes to Dinner: Arguments About the Ethics of Eating*, 268-294. New York: Routledge.

———. 2017. "Early Death and Later Suffering." In E. Gamlund & C. Solberg, eds. *Saving Lives from the Badness of Death,* Chapter 7. New York: Oxford University Press.

Moore, G.E. *Principia Ethica*. 1988. Amherst, New York: Prometheus Books.

Nagel, T. 1970. "Death." *Noûs* 4, 1: 73-80.

Narveson, J. 2012. "Future People and Us." In R. I Sikora & B. Barry, eds. *Obligations to Future Generations,* 38-60. Cambridge, UK: The White Horse Press.

Parfit, D. 1978. "Innumerate Ethics." *Philosophy & Public Affairs* 7, 4: 285-301.

———. 1984. *Reasons and Persons*. Oxford: Oxford University Press.

———. 2011. *On What Matters: Volume Two.* Oxford: Oxford University Press.

———. 2018. "Future People, the Non-Identity Problem, and Person-Affecting Principles." *Philosophy & Public Affairs* 45, 2: 118-157.

Roberts, M. 2003. “Can it Ever Be Better Never to Have Existed At All? Person-Based Consequentialism and a New Repugnant Conclusion.” *Journal of Applied Philosophy* 20, 2: 159-185.

Rosenbaum, S. 1986. “How to be Dead and Not Care: A Defence of Epicurus.” *Philosophical Quarterly*: 217-225

Scheffler, S. 2013. *Death and the Afterlife*. Oxford: Oxford University Press.

Singer, P. 2011. *Practical Ethics*. Cambridge: Cambridge University Press.

Smuts, A. 2012. “Less Good but Not Bad: In Defense of Epicureanism About Death.” *Pacific Philosophical Quarterly* 93: 197-227.

Suits, D. 2001. “Why Death Is Not Bad for the One Who Died.” *American Philosophical Quarterly* 38: 69–84.

Taylor, J.S. 2012. *Death, Posthumous Harm, and Bioethics*. New York: Routledge.

Temkin, L. 2000. “Equality, priority, and the leveling down objection.” In M. Clayton & A. Williams, eds. *The Ideal of Equality*. London: MacMillan

Temkin, L. 2012. *Rethinking the Good*. Oxford: Oxford University Press.

Timmerman, T. 2016. “Reconsidering Categorical Desire Views.” In M. Cholbi, ed. *Immortality and the Philosophy of Death*, 21-39. London: Rowman & Littlefield.

Uniacke, S. 1997. “Replaceability and Infanticide.” *The Journal of Value Inquiry*, 31: 153–166.

Velleman, D. 2015. *Beyond Price: Essays on Birth and Death.* Cambridge, UK: Open Book Publishers.

Weinberg, J. 2014. “Non-Identity Matters, Sometimes.” *Utilitas* 26, 1: 23-33.