

Extension and Replacement

Michal Masny
21 March 2023
mmasny@princeton.edu

Abstract: Many people believe that it is better to extend the length of a happy life than to create a new happy life, even if the total welfare is the same in both cases. Despite the popularity of this view, one would be hard-pressed to find a fully compelling justification for it in the literature. This paper develops a novel account of why and when extension is better than replacement which applies not just to persons, but also to non-human animals and humanity as a whole.

Word count: 9,100

1. Introduction

As Gustaf Arrhenius (2008, p. 211) observes,

It seems to be a widespread opinion that increasing the length of existing happy lives is better than creating new happy lives although the total welfare is the same in both cases, and that it may be better even when the total welfare is lower in the outcome with extended life.

Despite the popularity of this view, one would be hard-pressed to find a fully compelling justification for it in the literature. Arrhenius himself considers three candidates, but eventually rejects all of them. In this essay, I want to propose a novel account of why and when extension is better than replacement which appeals to a broadly conservative view about value.

Before I review a handful of rival justifications and introduce the details of my account, I would like to clarify three aspects of my discussion. First, I will understand the claim that extension is *better* than replacement in the wider, reason-implicating sense rather than in the narrower, axiological sense. That is, I take the relevant sentiment to be that we have a weighty moral reason to increase the length of existing happy lives rather than create new happy lives, although the total welfare is the same in both cases. This may be because the former state of affairs has overall

more value than the latter, but it need not be. Indeed, as I will argue, we often have a moral reason to choose extension over replacement, even if the latter would result in a more valuable state of affairs.

Second, while Arrhenius focuses on persons, I want to include other kinds of entities in my discussion. It seems to be an equally common view that extending *humanity's* tenure would be better than allowing another intelligent and sentient species to replace us. Likewise, we should reflect on whether extending the life of a happy *non-human animal* (hereafter just 'animal') is better than replacing it with another happy animal. (Opinions on this latter issue are admittedly more divided.) To my knowledge, these three questions have not been discussed together. This strikes me as a missed opportunity. As I propose here, there is a common ground for an important class of our moral reasons to favour extension over replacement with respect to animals, persons, and humanity. I will first derive this ground in the context of humanity as a whole, where the case for extension is arguably the most intuitive, and then expand my account to cover persons and animals.

Third, my aim in this paper is to identify *one important source* of our moral reasons to choose extension over replacement. It is not, however, to identify a *unique source* of such reasons. I choose to restrict my aim in this way mainly because concrete cases depicting the choice between extension and replacement typically have multiple features that could potentially ground a moral reason for extension. We could, of course, try to rule out many of these features by constructing increasingly intricate vignettes, but this approach runs the risk of foregoing tractability. Instead, I propose that we take the presented cases to illustrate the kinds of situations to which the proposed account applies, assess how common or rare these situations are, and in that way, gauge this account's explanatory power.

With these clarifications in place, my central question takes the following form:

Extension vs Replacement: What, if anything, grounds our moral reasons to choose extension over replacement with respect to non-human animals, persons, and humanity as a whole, when the total welfare is the same in both cases?

The answer I propose appeals to a broadly conservative view about value. This view emphasises the moral importance of preserving certain non-instrumentally valuable elements of humanity's culture (such as beautiful artworks,

cherished traditions, just institutions, bodies of scientific knowledge, and relations of equality) and of individuals' lives (such as loving relationships and important personal projects). In essence, my account holds that we have a moral reason to choose the extension of the life of an animal, person, or humanity over its replacement when and because this is necessary and sufficient for the preservation of those non-instrumentally valuable things.

Let me say three things about this view, in an effort to anticipate some natural questions. First, the proposed account implies that we *typically* have a moral reason to choose extension over replacement with respect to persons and humanity, but it does not imply that we *always* do. This, I believe, aligns with our intuitions. This account also goes some way towards explaining why people's opinions about the same choice regarding animals are divided. As I will explain in Section 4, it turns out that whether we have a conservative reason to extend animal lives depends on a largely empirical premise (namely, whether animal lives feature loving relationships and important personal projects which warrant preservation), and folk intuitions about this premise are divided.

Second, my answer has a *reductive* character. In my view, our conservative reasons to favour extension over replacement are grounded in the importance of preserving the valuable contents of an individual's life or the valuable elements of humanity's culture. They are not grounded in the importance of preserving life or the human species themselves, as a *non-reductive* account would have it. As I will explain in Section 5, while perhaps strange at first blush, the reductive account has two main advantages. For one thing, it's more parsimonious: to explain why we typically have a moral reason to choose extension over replacement, if we already believe that certain prudential goods and cultural artefacts are worth preserving, there is simply no need to postulate that, in addition, lives or the human species themselves are worth preserving. For another, this account better captures our considered judgments. I submit that there is no conservative reason to favour extension over replacement when all of the valuable elements of an individual's life or humanity's culture would be eradicated. In a slogan, what matters is the *continuity* of life rather than its mere *continuation*.

Finally, I would like to review a handful of views in population ethics which might seem to provide straightforward explanations of when and why extension is

better than replacement. Since many of the criticisms that I want to mention have been discussed extensively in the literature, I will be very brief.¹

Consider, first, two simple impersonal views. According to the average view, one outcome is better than another just in case it features higher average welfare. According to the critical level view, one outcome is better than another just in case it features higher total contributive welfare, where each individual's contributive value equals their welfare minus some positive constant representing the critical level. Both of these views imply that it is better to extend the life of a person or an animal rather than replace them with another individual because they prefer welfare to be spread among as few individuals as possible.

Alternatively, we might want to appeal to a person-affecting view. There are several ways in which such views can be spelled out—in terms of conditional goodness, harm-minimisation, or an asymmetry between comparative and non-comparative benefits and harms—but the animating idea appears to be that it's better to extend the life of an individual rather than allow for their replacement because the former option is better for someone and worse for no one.²

I think that none of these views provides a fully compelling justification of the intuition that extension is better than replacement. First, they all have well-known troubling implications for other issues in population ethics. For example, the average view and the critical level view entail the sadistic conclusion that it can be better to add to a population individuals with negative welfare rather than individuals with positive welfare. Moreover, person-affecting views struggle to explain why we have a moral reason not to create individuals with bad lives.³ Second, and this is more important in the present context, these views cannot provide a *unified* answer to my question. This is because they fail to explain why it is better to extend the tenure of humanity rather than allow its replacement with another intelligent and sentient species in cases when the continued existence of any particular individual is not at stake. Third, even if one of these views proves to correctly identify *a* moral reason to choose extension over replacement, this might not be the only or the most important reason that we have in this context. Indeed, the ensuing discussion seeks to show that an important class of our moral reasons to favour extension over replacement has been largely overlooked.

¹ For overviews of these views and criticisms, see Arrhenius (2008) and Greaves (2017).

² See, for example, McMahan (2013) and Frick (2020).

³ Though see Frick (2020) and Bader (2022) for recent attempts to accommodate this intuition within the person-affecting framework.

2. Humanity

What, if anything, grounds our moral reason to choose extension over replacement with respect to *humanity* as a whole when the total welfare is the same in both cases?

To get a better grip on this question, consider the following example. Suppose that you have two options. If you do X, humanity will continue to exist for another 500,000 years. If you instead do Y, humanity will go extinct, but another intelligent and sentient species will take its place and remain in existence for 500,000 years. Assume that these two futures would not differ in terms of total welfare, that doing Y does not amount to killing anyone, and that no one would die prematurely. Perhaps humans have recently become infertile, and you can either remedy this or create conditions under which another intelligent and sentient species would evolve on or settle our planet once all traces of our civilisation have disappeared. What, if anything, grounds your moral reason to do X rather than Y in this situation?

This question has received little philosophical attention in the literature.⁴ But there is another, related question which has been discussed more widely:

Extension and Extinction. What, if anything, grounds our moral reason to extend the length of humanity's tenure rather than to let it go extinct unreplaced?⁵

For example, in his recent book *Why Worry About Future Generations?* (2018), Samuel Scheffler argues that among the reasons to prevent the extinction of humanity are what we can call 'conservative reasons'. He writes:

All of the many things we value that consist in or depend on forms of human activity will be lost when human beings become extinct. No more beautiful singing or graceful dancing or intimate friendship or warm family celebrations or hilarious jokes or gestures of kindness or displays of solidarity. Other things that we value—physical artifacts, for example—may survive for a while, but with no one to appreciate their value, for in addition to the disappearance of *valuable things*, the extinction of the human race will mean the disappearance of *valuing* from the Earth. (pp. 69-70)

⁴ For one exception, see Williams (2006).

⁵ Henceforth, 'extinction' will refer to 'extinction without replacement'.

In this passage, Scheffler sketches an ingenious account of one important class of moral reasons to favour the extension of humanity's tenure over its extinction. It can be reconstructed as follows. We have a moral reason to favour extension over extinction because humanity's survival is necessary for the continued existence of valuing and certain valuable elements of humanity's culture, and we have an antecedent moral reason to preserve these things.

It is worth noting at the outset that while Scheffler's remarks focus on the *necessity* of the survival of humanity for the preservation of valuable things and valuing, an additional element is needed. Specifically, in any given set of circumstances, the survival of humanity must also be *sufficient* for preserving valuing and some valuable elements of humanity's culture are preserved. If the disappearance of valuing and the valuable elements of our culture was inevitable, there would be nothing to ground a conservative reason to prevent the extinction of humanity.

With this small amendment in place, let's consider: can the same kinds of considerations be used to support our moral reason to choose the extension of humanity's tenure over the replacement of humanity with another intelligent and sentient species?

There are two main grounds for hesitation. To begin with, consider valuing. On Scheffler's own view developed elsewhere, valuing X involves four elements: (i) a belief that X is valuable; (ii) a susceptibility to experience a range of context-dependent emotions regarding X; (iii) a disposition to experience these emotions as being appropriate; and (iv) a disposition to treat certain X-related considerations as reasons for action in relevant deliberative contexts (2010, p. 29). Based on this description, there seems to be nothing distinctively human about valuing that could not, or would not, be replicated in another intelligent and sentient species. Consequently, it looks like our concern for the continued existence of valuing cannot support a concern for the continued existence of humanity in particular, as opposed to the continued existence of intelligent and sentient life in general.

What about the valuable elements of humanity's culture, such as beautiful singing, intimate friendship, and displays of solidarity? One might likewise suggest that these things would not disappear altogether if humanity was replaced by another intelligent and sentient species. There is no principled reason to think that members of that species would not be capable of creating and performing music, establishing intimate connections with one another, and displaying virtues such as solidarity.

Thus, it may look like our concern for the continued existence of the valuable elements of humanity's culture cannot support a moral reason to extend humanity's tenure when replacement is the alternative.

But this second strand of scepticism rests on a misunderstanding of the conservative view about value that underlies Scheffler's account of our moral reasons to prevent extinction. On this view—also espoused by the likes of G. A. Cohen (2012), Johann Frick (2017), and Jacob Nebel (2015)—we have a moral reason to preserve certain valuable things even when they could be replaced by similar things of equal or greater value. Crucially, what we have a reason to preserve is not the existence of valuable things *in general*, or even certain *types* of valuable things, but rather *tokens* of valuable things. For example, we have a reason to preserve the iconic Golden Gate Bridge even if it could be replaced by another, even more impressive structure. This is important for thinking about the choice between extension and replacement. If another intelligent and sentient species were to replace us once all traces of our civilisation have disappeared, as we are asked to imagine, tokens of many valuable things present nowadays—*human* forms of singing, friendship, and solidarity—would not be preserved. Therefore, our concern for the continued existence of such things can, in fact, support a moral reason for extending humanity's tenure when the alternative is replacement with another intelligent and sentient species.

With these considerations in place, the conservative answer to my central question, as it concerns humanity as a whole, takes the following shape. We have a moral reason to extend the length of humanity's existence rather than allow its replacement with another intelligent and sentient species when and because this is necessary and sufficient for preserving certain valuable elements of our culture. These include distinctively human forms of singing, friendship, or solidarity, but also tokens of other non-instrumentally valuable things, such as beautiful artworks, sublime landscapes, unique languages, or relations of equality between members of the moral community.

Let's take a moment to emphasise two features of this view. First, conservative reasons are *pro tanto* reasons and must be weighed against other relevant moral considerations. For example, if the replacement of humanity with another intelligent and sentient species would eventuate in a state of affairs that would be vastly more valuable, then this might be what we should all-things-considered choose.

Second, conservative reasons have *varying strengths*. This strength is plausibly a function of both the quantity and the value of the things that are such that humanity's survival is necessary and sufficient for their preservation. One corollary of this is that our conservative reason to prevent extinction will likely be stronger than our conservative reason to prevent replacement. This is because, as I have suggested above, at least one of the things that are worth preserving—valuing itself—would disappear in the case of extinction but would not (or at least need not) disappear in the case of replacement.

Return to the main thread of the argument. The preceding discussion explained *why* we sometimes have a moral reason to choose extension over replacement. Now, let's examine *when* that's the case. To that end, it is helpful to single out three conditions which must be satisfied for such reasons to arise: *presence* (of non-instrumentally valuable things which warrant further preservation), *necessity*, and *sufficiency*. I think that these conditions are satisfied in most practically-relevant cases, but there are some noteworthy exceptions. Let me mention one example for each condition.

Take 'presence' first. As things stand, there is an abundance of non-instrumentally valuable things that warrant further preservation: thousands of beautiful artworks, sublime landscapes, unique languages, important traditions, collective endeavours, and other cultural artefacts. While some of them inevitably perish every year, the total number of these things keeps increasing, as a seemingly necessary by-product of human activity. But one day, all of this might change. A global nuclear conflict or a natural disaster of an unprecedented scale could reduce our planet and the fruits of our civilisation to rubble and dust.⁶ If humanity were to survive a catastrophe of this kind and find itself completely deprived of valuable things that warrant further preservation, there would be no conservative reason to choose extension over replacement (or extinction, for that matter). That is, even if we were sufficiently prepared (or lucky!) and had the opportunity to build a new world on top of the ruins of the old one, in the absence of genuine continuity of the human civilisation, extension would not be preferable to replacement from the conservative point of view.

Turn to 'necessity' now. For all we know, humanity is alone in our little corner of the universe. No one will rescue us if we get into trouble, and no one will

⁶ For recent illuminating discussions of how civilisational collapse could come about and what we could do to prevent it, see Ord (2020) and MacAskill (2022).

take care of our business once we're gone. But suppose, for the moment, that were not true. Perhaps an intelligent and sentient alien species has been observing us from afar and developed an appreciation for our ways of life. As long as things are going sufficiently well for us, these aliens will not interfere. But in the event of human extinction, their emissaries will rush to Earth and seamlessly take custody of as many valuable things as possible. They will safeguard our art, learn and practice our languages and traditions, and work hard to advance our collective endeavours. If that were to happen and nothing of value would be lost in this transition, there would be no conservative reason for extension over replacement either.⁷

Finally, consider 'sufficiency'. We can often expect humanity's survival to be sufficient for the preservation of many valuable elements of our culture. Some valuable elements of our culture are straightforward to preserve. Languages and traditions, for example, just need to be practised by people, and no special incentives are usually necessary for them to do that. In other cases, we already have elaborate systems in place: museums and galleries that house our artworks and whole professions dedicated to maintaining them in good condition and on display for people to appreciate.

But there are outliers. Imagine a scenario in which humanity faces a global catastrophe, and the only way for our species to survive is to forego everything of value. Perhaps a small number of people could find shelter in an underground bunker, where they would live in primitive conditions while everything on our planet's surface would perish. If truly none of the valuable elements of our culture would remain in this world, then there would be no conservative reason to extend humanity's tenure rather than allow for our replacement. This is so even if the surviving population could at a later time, when the conditions on the surface have sufficiently improved, create completely new cultural artefacts.⁸

This concludes the first part of our discussion. We can see that while we typically have a conservative moral reason to choose extension over replacement, there are exceptions. This strikes me as intuitively correct. I believe that what we ultimately care about—at least from the conservative point of view—is not the preservation of the human species itself, but rather the preservation of the valuable

⁷ Even if a few valuable elements would be lost in the transition (as they would be, perhaps, if upon all philosophers' sudden disappearance, philosophically-minded economists, physicists, and sociologists took over our positions and research agendas), the conservative reason for extension over replacement would be quite weak.

⁸ I will return to this issue in Section 5.

elements of our civilisation. These two things usually go together, but they can also come apart. Our species can survive even if our civilisation crumbles, and (as the scenario involving alien custodians is meant to illustrate) the fruits of our civilisation might survive even if our species goes extinct.

It is also worth noting that whether there is a conservative reason to choose extension over replacement is largely independent of welfare considerations. In particular, we could have a conservative reason to choose extension even if replacement would result in a state of affairs that would be more valuable. In fact, we could have a conservative reason for extending humanity's tenure even if it would result in a state of affairs with negative total welfare (for example, if future people suffered from some painful condition, but nonetheless preserved many valuable elements of our culture). That said, it is important to keep in mind that conservative reasons are merely *pro tanto* reasons. Even if they support extension in such cases, they do not guarantee that this is what we should all-things-considered choose.

3. Persons

Turn now to the second part of my central question: what, if anything, grounds our moral reason to choose extension over replacement with respect to *persons* when the total welfare is the same in both cases?

This choice can be illustrated as follows. You have two options. If you do X, Ada will continue to exist for another 50 years. If you instead do Y, Ada will die, but Adam will come into existence and live for 50 years. Assume that Ada's extra 50 years and Adam's 50-years-long life would not differ in terms of welfare, and that doing Y does not amount to killing or failing to rescue Ada. Perhaps you must allocate one unit of some scarce drug, and that drug will eventually either extend Ada's life or else cause Adam to exist. What, if anything, grounds your moral reason to do X rather than Y in this situation?

I believe that, in this context as well, the answer has to do with the importance of preserving valuable things. In particular, we have a moral reason to extend a person's life rather than create another life when and because this is necessary and sufficient to preserve certain valuable things.

What are the relevant valuable things? On the one hand, there are things which are *good simpliciter*. These are the things we discussed in the previous section, such as beautiful artworks, important traditions, and relations of equality between

members of the moral community. In some cases, the continued existence of a particular person will be both necessary and sufficient for the preservation of such things. For example, we could imagine that preserving a magnificent but crumbling cathedral depends on the continued oversight of a renovation project by a particular architect with unique expertise. Likewise, the maintenance of peace in some region could require the continued tenure of a particular, extraordinarily skilful politician. We would have moral reasons to ensure the continued existence of these individuals, rooted in the concern for preserving valuable things. That said, such cases will be rare.

More commonly, our moral reasons to extend a person's life will be grounded in the importance of preserving *prudential goods*: things which are valuable for particular people. While discussions of conservatism about value in the literature have focused almost exclusively on things which are good simpliciter, I have argued in other work that we also have a distinctive moral reason to preserve certain prudential goods, such as loving relationships and important personal projects, even when a superior replacement is available.⁹

Two main examples motivate this idea. The first concerns Jim Brown, a former fullback for the Cleveland Browns. In 1966, Brown shocked the football world by announcing his retirement from the sport at the young age of 29. His decision was not motivated by injury or a decline in performance. In fact, he had just won the NFL Most Valuable Player award for the third time. Instead, Brown left football to pursue another career, acting, where he also enjoyed great success. He became the first Black action movie star, appearing in over 50 movies and TV shows, many of which were blockbusters.

The second example is fictional. Suppose that Ann is in a loving relationship with Beth. One day, Ann meets Sally, who happens to be very similar to Beth in all respects except one. Like Ann, Sally is a dedicated runner, so they could enjoy their training together and support each other during races, bringing them even closer. After some deliberation, Ann leaves the relationship with Beth for Sally.

What should we think about these cases? If you are like most people, Brown's decision seems, in one respect, regrettable. Crucially, this is not just because he missed out on the many athletic accomplishments that lay ahead of him, but also because of what he had left behind. He quit a career that was a non-instrumentally valuable personal project that he gave up much for and that still had plenty to give

⁹ See Chapter 3 of this dissertation.

back to him. Likewise, there is something troubling about Ann's decision. That old relationship was non-instrumentally valuable and would have continued to be so, for all we know. In other terms, despite the presence of an equally good or even better replacement, both choices leave us with considerable normative residue. This normative residue, I believe, indicates that we have a distinctive moral reason to conserve important personal projects and loving relationships—just as we have a distinctive moral reason to conserve beautiful paintings, important traditions, or relations of equality.¹⁰

If this is right, then the conservative account can potentially extend to a much wider range of cases than we might have originally assumed. Few people are indispensable for the preservation of things which are valuable simpliciter, such as architectural marvels or peaceful relations in any given region. But almost everyone has some loving relationships or important personal projects, and everyone who does is indispensable for their preservation. Thus, it's almost always the case that we have a conservative reason to choose extension over replacement with respect to persons.

There are exceptions, however. Each of the three conditions—presence, necessity, and sufficiency—can be frustrated. I will give one example for each.

First, in some instances, the kinds of goods which warrant preservation were once present in a person's life but no longer are. This is sometimes true of elderly people: they could have already lost all their loved ones, either completed or abandoned their important personal projects, and severed their connections with valuable cultural artefacts. In such cases, there is no relevant ground for a conservative reason to choose extension over replacement.

Second, turn to the necessity condition. As suggested earlier, this condition is trivially satisfied with respect to prudential goods. The fruits of human civilisation could perhaps survive without the human species, but one's relationships and personal projects could not survive without them. However, the necessity condition is not trivially satisfied with respect to things which are good simpliciter. Consider again the example of an architect who oversees the efforts to preserve a magnificent but crumbling cathedral. If there is another expert who would seamlessly take over this role in the event of the first architect's death, then there is no corresponding ground for a conservative reason to choose extension over replacement either.

¹⁰ It is worth emphasising that this reason is of moral, not prudential, variety. After all, Brown had a reason to preserve his first career, and Ann had a reason to preserve her first relationship even though their new engagements would have been at least as good for them.

Finally, consider sufficiency. We have previously discussed a scenario in which humanity can only survive in some diminished state or at the expense of the fruits of our civilisation. There is a familiar analogue in the present context. When people describe their experience with chemotherapy or other aggressive forms of cancer treatment, they often speak about fatigue, brain fog, mood changes, stress, and loss of sexual function. These side effects might seem relatively insignificant when considered in isolation, but they tend to carry further deleterious consequences, especially when they occur in combination. For instance, they might prevent a person from engaging in fulfilling social interactions, valuing their professional and athletic goals, or maintaining serious intellectual engagements. That is, as a result of undergoing these forms of treatment, a person could lose all of their loving relationships and important personal projects. If that were the case, there would be no conservative reason to choose extension over replacement.

To summarise this part of the discussion, the conservative account implies that we *almost always* have a conservative moral reason to extend a person's life rather than to create a new life. The only exceptions are cases in which nothing of value would be preserved alongside the person's life.

Crucially, these conservative reasons are *agent-neutral*. We have them regardless of whether we are personally involved in the relationships or important projects of the person whose life we are in a position to extend. In this respect, preserving a life which features loving relationships and important personal projects is like preserving the Grand Canyon or the Mona Lisa.

4. Animals

Turn now to the third and last entity on our list, non-human animals. We can consider a similar choice in this context. You have two options. If you do X, a cow named Moo will continue to exist for another 5 years. If you instead do Y, Moo will die, but another cow named MooTwo will come into existence and live for 5 years. Assume that Moo's extra 5 years and MooTwo's 5-years-long life would not differ in terms of total welfare. Suppose also that doing Y does not amount to killing or failing to rescue Moo. Perhaps you must allocate one unit of a scarce drug, and that drug will either eventually extend Moo's life or else cause MooTwo to exist. What, if anything, grounds your moral reason to do X rather than Y in this situation?

In my experience, people's intuitions about this case are more divided than about the previously considered cases. Some people, including me, believe that there

is a moral reason to extend the lives of at least some animals—maybe not flies or shrimp, but certainly cows, dogs, and chimpanzees. Others, including many of those who accept the practice of ‘humane omnivorism’, believe otherwise.¹¹

What does the conservative account say about this issue? As before, we have a moral reason to choose the extension of an animal’s life over its replacement with another animal when and because this is necessary and sufficient to preserve certain valuable things. Needless to say, this is a conditional claim, so we need to examine when, if ever, the antecedent holds true.

Start with things which are valuable simpliciter, such as beautiful artworks or relations of equality between people. There are some possible cases in which a particular animal’s continued existence is essential to the continued existence of these things. Perhaps if a dog named Fido dies, then its architect-owner will fall into the grips of depression, and the magnificent cathedral whose renovation he has been overseeing will crumble. But this kind of dependence is of limited philosophical interest in the present context and such cases would be extremely rare anyway.

What about prudential goods? Take loving relationships. Empirical research on animal behaviour strongly suggests that many animals develop lasting ties with each other (for example, orangutan mothers stay with their offspring for eight to ten years), experience emotions characteristic of loving relationships (such as grief), and are willing to make sacrifices when others of their own kind are at risk of harm.¹²

The issue, however, is whether these bonds count as ‘genuine loving relationships’. Presumably, only these kinds of relationships warrant preservation. By contrast, mere acquaintances and fleeting connections do not seem to have this profile.

According to one popular view, defended by Niko Kolodny (2003, p. 150),

Love is a kind of valuing. Valuing X, in general, involves (i) being vulnerable to experience certain emotions regarding X, and (ii) believing that one has reasons both for this vulnerability to X and for actions regarding X.

¹¹ Humane omnivorism is the practice of raising animals in humane conditions, killing them for consumption, and replacing them with individuals, which have lives that are at least as good. For overviews of arguments in favour and against this practice, see McMahan (2008) and Delon (2016).

¹² For an informative overview of some of this empirical work, see Gruen (2021).

Some animals satisfy the first condition, as evidenced by their displays of emotions such as grief. What's more controversial is whether they satisfy the second condition: that is, whether they see their relationships as reason-giving. Although, as noted earlier, there does not seem to be anything distinctively human about valuing, certain sophisticated cognitive capacities do seem to be required to recognise one's relationship as a source of reasons. Kolodny himself suggests that small children likely lack the relevant capacities and therefore do not stand in genuine loving relationships with their parents (2003, fn. 22). If that's right, then animals presumably also lack these capacities and thus fail to have genuine loving relationships.

Turn now to personal projects. Empirical research has shown that many animals engage in temporally extended and goal-oriented activities. To give just one example, certain species of birds (such as scrub jays) cache store food to recover it hours, days, or even weeks later when hungry, which suggests that they can anticipate and plan for future contingencies.¹³

But do these activities count as personal projects of the kind that is worth preserving? According to a popular view, genuine relationships and projects are similar in that they must be valued by an individual. For example, Scheffler (2010, p. 48) writes that

Valuing a personal project, like valuing a personal relationship, involves seeing it as reason-giving. In other words, to value a project of one's own is, among other things, to see it as giving reasons for action in a way that other people's projects do not, and in a way that other comparably valuable activities in which one might engage do not.

Overall, these considerations reveal that whether we have a conservative reason to extend rather than replace animal lives hinges on a largely empirical matter: whether animals have the kinds of cognitive capacities that are required for genuine loving relationships and personal projects. I will not attempt to settle this question here because it requires expertise in animal studies that I do not have. Instead, I will end this section by emphasising two insights from the preceding discussion.

¹³ See Correia et al. (2007) for a recent study and Dickinson (2011) for an overview of the literature on goal-directed behaviour and future planning in animals.

The first insight is that those who find the practice of humane omnivorism morally troubling, as I do, now have a new line of critique at their disposal. They can argue that we have a moral reason to extend rather than replace animal lives when and because these animals' continued existence is necessary and sufficient for preserving certain valuable things like these animals' loving relationships and important personal projects. 'All' they have to do is show that these valuable things are, in fact, present in animal lives.

Of course, this might not be enough to conclude that humane omnivorism is always impermissible. Conservative reasons are merely pro tanto reasons, so they must be weighed against moral considerations that arguably support the practice of humane omnivorism, such as those concerning the pleasure some people derive from meat consumption. But I believe that conservative reasons are strong enough to make it difficult to outweigh them in this way.

Moreover, one attractive feature of this line of argument is that conservative reasons apply to a wide range of ways in which humane omnivorism could be practised. For example, they apply not only to an arrangement in which animals are killed prematurely but also to a situation in which animals are genetically engineered to drop dead when in prime condition for consuming their flesh. From the conservative point of view, what matters is not whether the animal would have survived otherwise (had we not consumed it), but whether it is in our power to ensure that it stays alive. In particular, if a pill or some other form of treatment could mitigate the effects of genetic engineering and secure additional years of life for the animal, we would have a pro tanto conservative reason to provide it.

The second insight is of sociological nature. I noted earlier that people's intuitions are divided with respect to whether we have a moral reason to choose extension over replacement for animals. The conservative account suggests that this normative disagreement might have its root in a disagreement about a largely empirical matter: whether animal lives feature genuine, non-instrumentally valuable loving relationships and important personal projects.

5. Continuation and continuity

As should be clear by now, the account I have proposed has a reductive character. Our conservative reasons to extend the life of a person or an animal are grounded in the importance of preserving the valuable contents of their life, such as their loving relationships and important personal projects. Likewise, our reasons for

extending humanity's tenure are grounded in the importance of preserving the valuable elements of humanity's culture, such as important traditions, beautiful artworks, and unique languages. On this view, what we ultimately care about from the conservative point of view is the *continuity* of an individual's life or humanity's tenure rather than the mere *continuation* of their existence.

Some people who are sympathetic to conservatism about value might find this approach strange. In particular, they might insist that, in addition to preserving the valuable contents of an individual's life and the valuable elements of humanity's culture, conservatives about value should care about preserving the individual's life and humanity (understood as the human species) in themselves. To deny that, the thought goes, would be to unduly instrumentalise the value of the individual's life and humanity. We can call this *the non-reductive view*.

One proponent of the non-reductive view is Johann Frick (2017, p. 359), who asserts the following in his discussion of the reasons to prevent the extinction of humanity:

It is commonplace to claim of a wide range of things that they have final value in this sense: wonders of nature, great works of art, animal and plant species, languages, culture, etc. The suggestion that *humanity* too, with its unique capacities for complex language use and rational thought, its sensitivity to moral reasons, its ability to produce and appreciate art, music, and scientific knowledge, its sense of history, and so on, should be deemed to possess final value, therefore strikes me as extremely plausible. I do not, however, have the space to argue this claim in this article. I will ask you to grant it to me as a premise, in the interest of seeing whether the final value of humanity may ground a moral reason to ensure humanity's survival. What I shall argue in the following is that there is a link between responding *appropriately* to the final value of humanity and being at least disposed to ensure its survival. (p. 359)

In the same spirit, one might think that individual humans (and perhaps animals) have final value in virtue of their unique rational and emotional capacities, and that we have conservative reasons to preserve them that go above and beyond our reasons to preserve the valuable contents of their lives. On this view, the mere continuation of an individual's or humanity's existence is within the scope of the conservative concern (as long as the relevant capacities are retained).

I think that we should reject the non-reductive view for two main reasons. The first reason has to do with theoretical parsimony. As evidenced by the quote above, proponents of the non-reductive view accept that the valuable elements of our culture (and presumably the valuable contents of our lives) are in themselves worth preserving. But, as my argument shows, that's all we need to explain the intuition that we almost always have a weighty moral reason to choose extension over replacement. There is simply no need to postulate that individual animal lives, individual human lives, and humanity are *also* worth preserving in themselves. Thus, on the grounds of theoretical parsimony, we should reject the non-reductive view in favour of the reductive view.

The second reason has to do with extensional adequacy. I believe that the non-reductive view over-generates: there are cases in which we intuitively lack a conservative reason to choose extension over replacement, but the non-reductive view implies that we do.

To be sure, the non-reductive view need not imply that we always have a moral reason to extend the life of an animal, person, or humanity. For example, in the previously mentioned paper, Frick (2017) suggests that we would not have a conservative reason to ensure the survival of humanity if it were to irreversibly lose its unique capacities, such as those for producing and engaging with valuable cultural artefacts. Presumably, proponents of the non-reductive view would also say that we would have no conservative reason to extend the life of a person (or an animal) if they were to lose the rational and emotional capacities that allow them to initiate and maintain loving relationships and important personal projects. In this respect, the non-reductive view and the reductive view are in agreement.

However, the non-reductive view implies that we would have a conservative reason to extend humanity's tenure rather than allow its replacement with another intelligent and sentient species even if every single valuable element of humanity's culture would be eradicated: all of our traditions, languages, works of art, technologies, written and oral histories, and ways of life. Likewise, the non-reductive view implies that we would have a conservative reason to extend the life of a person (or an animal) rather than allow their replacement with another person even if every single valuable element of their life would perish, including all of their loving relationships and important personal projects. All that appears to matter to the non-reductivist in these cases is a brute fact about numerical identity: that the existence of a particular individual or particular species is extended.

In this respect, I believe, the non-reductive view betrays the conservative ethos. Take the example of a beautiful painting. Suppose that we can either erase it and create another beautiful artwork on the same canvas, or else destroy the first canvas entirely and create another artwork on a different canvas. There is no conservative reason to do the former rather than the latter. I submit that the mere continued existence of a particular individual or species is nothing more than the continued existence of a canvas. We have no conservative reason to prefer the continued existence of an individual whose life has two discontinuous halves and features entirely different contents (effectively, an individual with two distinct lives) over the consecutive existence of two individuals with corresponding lives. Likewise, we have no conservative reason to prefer the continued existence of a single (human) species whose history has two discontinuous halves and features entirely different cultures (effectively, a species with two distinct civilisations) over the consecutive existence of two species with corresponding cultures. This, I think, gives us a decisive reason to reject the non-reductive view.¹⁴

6. Further issues

The first part of this paper identified a common ground for our moral reasons to favour extension over replacement with respect to humanity, persons, and animals. In this section, I want to address two further issues connected to that discussion.

The first issue concerns the methods that could be used to extend a person's (or an animal's) life. In recent years, we have made significant progress in our

¹⁴ Suppose, for the moment, that I am mistaken, and we have good reasons to accept the non-reductive view. Does that invalidate my project of developing the reductive view? I don't think so. Strictly speaking, the non-reductive view and the reductive view are compatible with each other. If that's the case, then in typical cases, we would have a conservative reason to extend humanity's tenure in virtue of the fact that the human species is itself worth preserving *and* a conservative reason to do that in virtue of the fact that its survival is necessary and sufficient for the continued existence of certain valuable elements of our culture which are in themselves worth preserving. This would be just another instance of the commonplace ethical phenomenon of *normative overdetermination*. Crucially, overdetermination does not mean redundancy here. If there are both non-reductive and reductive conservative reasons to choose extension over replacement, then the conservative support for extension is in one way stronger. This is important because, as I noted earlier, conservative reasons are pro tanto reasons and have to be weighed against other moral considerations, such as the prospective value of the replacement outcome. So, perhaps, we should choose extension over replacement even if the total welfare in the latter scenario is much higher, and not just slightly higher.

scientific understanding of nutritional, lifestyle, pharmacological, and genetic interventions that can potentially slow down or even reverse ageing.¹⁵ However, I want to focus on two more speculative routes to a longer life that continue to feature prominently in the popular imagination: cryonics and uploading. Cryonics is the process of storing a person's brain or whole body at a very low temperature after their legal death in the hope that scientific progress will one day make revival possible. Uploading, on the other hand, involves scanning a person's brain and uploading this information to a computer, where their brain is then simulated.

Concerns about feasibility aside, philosophical discussions of these potential methods of life extensions have focused on two issues. The first is whether these methods would preserve a person's consciousness.¹⁶ The second is whether they would preserve one's memories and thus their psychological continuity, which many philosophers regard as necessary for personal identity.¹⁷

My conservative account puts a spotlight on a third, novel kind of worry. In a nutshell, even if these life extension methods were consciousness- and identity-preserving, something else could be lost in the process: one's loving relationships and important personal projects.

To see that, note that our loving relationships and important personal projects are embedded in the external world. This is important in a number of respects. Most obviously, our loving relationships are relationships with other people. If a person undergoes cryonics and is revived decades later in a world from which their loved ones are absent, their loving relationships will not survive either. Likewise, if a person is uploaded, but their loved ones stay in the physical world, it will be impossible, or at least more difficult, to maintain these loving relationships. Similar considerations apply to our important personal projects. While they need not involve other people, they generally depend for their purpose and value on other features of the external world. For example, the project of being the first person to run the marathon under two hours ceases to make sense if that has already been done by someone else (as will likely be the case if you spend decades in a cryogenic tank) or if you no longer have a physical body. Thus, extending one's life in one of these ways would often involve sacrificing these prudential goods, even if one's

¹⁵ For overviews of recent developments in this field, see Kennedy et al. (2014), Cohen (2015), Newman et al. (2016), and Partridge et al. (2020).

¹⁶ See Chalmers (2010).

¹⁷ For this worry, see Chalmers (2010), Aaronson (2016), and Doyle (2018). For an overview of the literature on personal identity, see Olson (2022).

consciousness and identity were preserved. In such cases, we would lack a conservative reason to choose the extension of a person's life through cryonics or uploading over their replacement.

The second issue connected to my discussion of the choice between extension and replacement is this. We have previously considered cases, including the ones just above involving cryonics and uploading, in which extending the existence of a person or humanity is not sufficient for the preservation of any valuable things, and so there is no conservative reason to choose extension over replacement. But perhaps an even stronger conclusion is warranted in these scenarios: namely, that we have a conservative reason to choose replacement over extension.

To assess that, let's return to the case in which a person can undergo an aggressive form of cancer treatment that would make them suffer from severe fatigue, brain fog, mood changes, stress, and loss of sexual function. These side-effects, we have stipulated, would prevent a person from engaging in fulfilling social interactions, valuing their goals, and maintaining serious intellectual engagements. Suppose now that this person would not entirely lose their loving relationships and important personal projects, but they would instead survive in some significantly diminished state.

What does the conservative view about value say about this case? Plausibly, conservative reasons go beyond merely ensuring that valuable things continue to exist. For example, when someone lets an extraordinarily beautiful painting deteriorate so much that it survives at a level at which it is just barely attractive, their behaviour clearly falls short of being a fitting response to the painting's value. This suggests that it is important to preserve valuable things in a prosperous condition. Now, in the case of the cancer patient, the concern is precisely that, in virtue of staying alive, that person would fall short of this standard and thereby mistreat their relationships and projects. This suggests that there would actually be a conservative reason to choose replacement over extension in this case.

To be sure, if the patient does not undergo the treatment and dies soon, their loving relationships and important personal projects will come to an end. This might seem like an even worse fate for these valuable things compared to surviving in a diminished state. However, what matters for the conservative assessment of these scenarios are not just the *end states*, but also the *manner* in which they will come about. It is one thing to let an important text or a religious emblem tatter and find its way to a landfill, and another to bury or burn it in a ceremonial way. Likewise, it

seems more fitting to allow our loving relationships and important personal projects to end when they are still thriving than to have them survive in a scarred condition.¹⁸

Analogous considerations apply to humanity as a whole: in some circumstances, we may have a conservative reason to let it be replaced by another intelligent and sentient species rather than to extend its tenure. This would be so if ensuring the survival of humanity was possible only at the cost of having many valuable elements of our culture survive in a significantly diminished state. I won't repeat the entire argument for this conclusion.

Instead, I want to use this opportunity to connect the discussion of this issue with a famous claim made by Derek Parfit. In the closing section of *Reasons and Persons* (1984), Parfit asks us to compare three outcomes: (1) peace; (2) a nuclear war that kills 99% of the world's population; and (3) a nuclear war that kills everyone. He claims that while (2) is worse than (1), (3) is 'very much' worse than (2) because a nuclear war that kills everyone would foreclose human history, which might otherwise involve billions of extraordinarily happy lives and many spectacular achievements.

Parfit might be ultimately right here. But I think that it is important to recognise that, *in one respect*, the scenario in which humanity goes completely extinct could be better than the scenario in which a tiny fraction of the world's population survives. This would be so when and because the survival of humanity in this diminished state would degrade the valuable elements of our culture, such as important traditions, great works of art, just institutions, scientific knowledge, and relations of equality between members of the moral community. As in the case of beautiful paintings, loving relationships, and important personal projects, it seems more fitting to allow the valuable elements of our culture to disappear when they are still thriving than to have them survive in a scarred condition.

7. Concluding remarks

What, if anything, grounds our moral reasons to choose extension over replacement with respect to non-human animals, persons, and humanity as a whole?

¹⁸ Sometimes, one might have a third option: to extend their life but end all of their engagements before they perish. However, it is not possible to terminate in a dignified way all of one's valuable attachments with the intention of, say, experiencing a few extra years of pleasure. So, conservative reasons tell against this alternative.

According to the conservative account that I have presented, we have a weighty moral reason to choose extension over replacement when and because extension is necessary and sufficient for the preservation of certain valuable things, such as cultural artefacts, loving relationships, or important personal projects. This conservative account does not imply that we always have a moral reason to choose extension over replacement. However, in the context of persons and humanity as a whole, it implies that we almost always do. In the case of animals, on the other hand, the jury is still out: we must first determine whether their lives feature genuine loving relationships and important personal projects of the kind that warrants further preservation.

References

- Aaronson, S. (2016). Ghost in the Quantum Turing Machine. In S. B. Cooper and A. Hodges, eds. *The Once and Future Turing: Computing the World*. Cambridge: Cambridge University Press.
- Arrhenius, G. (2008). Life Extension Versus Replacement. *Journal of Applied Philosophy*, 25 (3): 211-227.
- Bader, R. (2022). The Asymmetry. In J. McMahan, T. Campbell, J. Goodrich, & K. Ramakrishnan (eds.). *Ethics and Existence: The Legacy of Derek Parfit*. New York: Oxford University Press.
- Chalmers, D. (2010). The Singularity: A Philosophical Analysis. *Journal of Consciousness Studies*, 17 (9-10): 7-65.
- Cohen, G. A. (2012). Rescuing Conservatism: A Defense of Existing Value (All Souls Version). In his *Finding Oneself in the Other*, edited by Michael Otsuka. Princeton: Princeton University Press.
- Cohen, J. (2015). Death-defying experiments: pushing the limits of life span in animals could someday help lengthen our own. *Science*, 350 (6265): 1186-1187.
- Correia, S., Dickinson, A., & Clayton, N. (2007). Western Scrub Jays Anticipate Future Needs Independently Of Their Current Motivational State. *Current Biology*, 17 (10): 856-861.
- Delon, N. (2016). The Replaceability Argument in the Ethics of Animal Husbandry. In: Thompson, P., Kaplan, D. (eds) *Encyclopedia of Food and Agricultural Ethics*. Springer, Dordrecht.
- Dickinson, A. (2011). Goal-Directed Behavior and Future Planning in Animals. In R. Menzel and J. Fischer, eds. *Animal Thinking: Contemporary Issues in Comparative Cognitive*. Cambridge: MIT Press.

- Doyle, D. J. (2018). *What Does it Mean to be Human? Life, Death, Personhood and the Transhumanist Movement*. Berlin: Springer.
- Frick, J. (2017). On the Survival of Humanity. *Canadian Journal of Philosophy*, 47(2-3): 344-367.
- Frick, J. (2020). Conditional Reasons and the Procreation Asymmetry. *Philosophical Perspectives*, 34: 53-87.
- Greaves, H. (2017). Population Axiology. *Philosophy Compass*, 12(11): 1-15.
- Gruen, L. (2021). The Moral Status of Animals. In: E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Summer 2021 Edition.
- Kennedy, B. K., Berger, S. L., Brunet, A., Campisi, J., Cuervo, A. M., Epel, E. S., Franceschi, C., Lithgow, G. J., Morimoto, R. I., Pessin, J. E., Rando, T. A., Richardson, A., Schadt, E. E., Wyss-Coray, T., & Sierra, F. (2014). Geroscience: linking aging to chronic disease. *Cell*, 159 (4): 709–713. 10.1016/j.cell.2014.10.039
- Kolodny, N. (2003). Love as Valuing a Relationship. *The Philosophical Review*, 112(2): 135-189.
- Masny, M. (ms). Conservatism about Prudential Value.
- McMahan, J. (2008). Eating Animals the Nice Way. *Daedalus*, Winter 2008: 1-11.
- McMahan, J. (2013). Causing People to Exist and Saving People's Lives. *The Journal of Ethics*, (17)5: 5-35.
- Nebel, J. (2015). Status Quo Bias, Rationality, and Conservatism about Value. *Ethics*, 125(1): 449-476.
- Nebel, J. (2022). Conservatisms about the Valuable. *Australasian Journal of Philosophy*, 100(1): 180-194.
- Newman, J. C., Milman, S., Hashmi, S. K., Austad, S. N., Kirkland, J. L., Halter, J. B., & Barzilai, N. (2016). Strategies and Challenges in Clinical Trials Targeting Human Aging. *The journals of gerontology. Series A, Biological sciences and medical sciences*, 71 (11), 1424–1434. <https://doi.org/10.1093/gerona/glw149>
- Olson, E. T. (2022). Personal Identity. In: E. N. Zalta, *The Stanford Encyclopedia of Philosophy*, Summer 2022 Edition.
- Parfit, D. (1984). *Reasons and Persons*. Oxford University Press.
- Partridge, L., Fuentealba, M., and Kennedy, B. K. (2020). The quest to slow ageing through drug discovery. *Nature Reviews Drug Discovery*, 19 (8): 513-532.
- Scheffler, S. (2010). *Equality and Tradition: Selected Essays*. New York: Oxford University Press.
- Scheffler, S. (2018). *Why Worry About Future Generations?* New York: Oxford University Press.