Philosophical Studies (2025) 182:1115–1132
https://doi.org/10.1007/s11098-025-02320-0

Extension and replacement

Michal Masny[1]

Accepted: 22 March 2025 / Published online: 4 April 2025

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 that applies not just to persons but also to non-human animals and humanity as a whole.

Keywords Value theory · Population ethics · Replacement · Life extension · Existential risk · Conservatism

1 Introduction

As 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 propose a novel account of why and when extension is better than replacement that appeals to a broadly conservative view about value.

✉ Michal Masny
mmasny@mit.edu

[1] Massachusetts Institute of Technology, Cambridge, United States

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-implying sense rather than in the narrower, axiological sense. That is, I take the relevant sentiment to be that we have a weighty pro tanto 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. (Intuitions regarding the 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 identify 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 and 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 that I propose appeals to a view known as *conservatism about value*. This view holds that we have a pro tanto moral reason to conserve certain non-instrumentally valuable elements of our culture (such as beautiful artworks, unique languages, cherished traditions, just institutions, and bodies of knowledge) and individuals' lives (such as loving relationships and important personal projects). At the minimum, this means that we have a moral reason not to destroy these things, even when they could be replaced by other things of equal or even greater value. For example, we seem to have a moral reason to preserve the marvellous Golden Gate Bridge, even if we could

build a more impressive structure in its place.[1] By contrast, we have no such reason to preserve an ordinary $10 bill when it could be replaced with another $10 bill. I take this position to be highly intuitive. In my discussion, I am going to assume that it's correct and instead focus on its ability to explain our intuitions about extension and replacement.[2]

In essence, my account holds that we have a moral reason to extend the life of an animal, person, or humanity rather than replace it when and because this is necessary and sufficient for the preservation of those non-instrumentally valuable elements of our culture or individuals' lives. 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 Sect. 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 that warrant preservation), and folk intuitions about this empirical 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 Sect. 5, while perhaps strange at first blush, the reductive account has two main advantages: it's more parsimonious and better captures our considered judgments.

Finally, I would like to critically review a handful of views in population ethics that might seem to provide straightforward explanations of when and why extension is better than replacement. Since many of the criticisms that I mention have been discussed extensively in the literature, I will be very brief.[3]

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 value, where each individual's contributive value equals their welfare minus some positive constant representing the critical level. Both

[1] In this and other cases, I set aside uncertainty regarding the outcome that is ubiquitous in the real world.

[2] For other discussions of conservatism about value, see Scheffler (2007, 2010, 2018), Cohen (2012), Nebel (2015, 2022), Frick (2017), and Masny (forthcoming). Some of these authors bring up issues related to my topic here. For example, Frick (2017) and Scheffler (2018) discuss human extinction. Nebel (2015, 2022) brings up the biblical story of Job, whose children are killed and then replaced by Cohen (2012) considers the possibility of humans turning into 'superhumans' through a series of biomedical enhancements. And in Masny (forthcoming), I discuss whether it's better if the history of humanity features a pattern of improvement rather than a pattern of deterioration, other things being equal. I will have more to say about Frick's and Scheffler's views shortly, but now I just want to emphasise that the choice between extension and replacement with respect to persons, animals, and humanity has not yet been systematically addressed in this literature.

[3] For overviews, see Arrhenius (2008) and Greaves (2017).

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 effectively prefer value to be spread among as few individuals as possible.

Alternatively, one 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, in the present context, the core 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.[4]

I think that none of these views provides a fully satisfying justification for 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 so-called 'sadistic conclusion' that it can be better to add to a population some number of individuals with negative welfare rather than a group of individuals with positive welfare.[5] Moreover, person-affecting views struggle to explain why we have a moral reason not to create individuals with bad lives.[6] Second, and this is more important in the present context, these views cannot provide a *unified* answer to my question. This is because they cannot 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. Finally, 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. A full and accurate normative description of the choice between extension and replacement might have to include conservative reasons anyway.

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 stylised 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 become extinct, but another intelligent and sentient species will take our 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 a situation such that another intelligent and sentient species will evolve or settle on 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?

[4] See, for example, McMahan (2013), McDermott (2019), and Frick (2020).

[5] Arrhenius (2000).

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

This question has received little philosophical attention in the literature.[7] But there is, of course, a 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 become extinct unreplaced?[8]

For example, in *Why Worry About Future Generations?* (2018), 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)

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, which 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 the latter.

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. If the disappearance of valuing and our culture were inevitable, there would be nothing to ground a conservative reason to prevent the extinction of humanity.[9]

With this small amendment in place, I would like to now consider whether these conservative considerations can also be used to support our moral reason to choose the extension of humanity's tenure over its replacement 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) susceptibility to experience a range of context-dependent emotions regarding X; (iii) disposition to experience these emotions as being appropriate; and (iv) disposition to treat certain X-related considerations as reasons

[7] For one exception, see Williams (2006).

[8] Henceforth, 'extinction' will refer to 'extinction without replacement'.

[9] I will discuss a case of this sort later in this section.

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, displaying virtues such as solidarity, and creating other valuable cultural artefacts. 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.

However, this second strand of scepticism rests on a misconception about conservatism about value that underlies Scheffler's account of our moral reasons to prevent extinction. On this view, we have moral reasons to preserve certain valuable things even when they could be replaced by similar things of equal or greater value. Crucially, what we have reasons 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 architectural 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—particular (human) forms of singing, friendship, and solidarity—would not be preserved. Therefore, our concern for the continued existence of valuable things does, in fact, support a moral reason for extending humanity's tenure when the alternative is replacement with another intelligent and sentient species. This is so even if the new (alien) forms of signing, friendship, and solidarity turned out to be more valuable than their (human) predecessors.[10]

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 human 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.

[10] What matters here first and foremost is that the presently existing human forms of signing, friendship, and solidarity are 'presently existing', not that they are 'human'. In a reverse case, members of an alien species would have a conservative moral reason to extend their tenure even if they would otherwise be replaced by a human species, whose corresponding cultural artefacts would be even more valuable.

Let's take a moment to emphasise two features of this answer. 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 vastly more valuable state of affairs, then this might be what we should, all things considered, choose.

Second, conservative reasons have *varying strengths*. Their 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 would 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.

The preceding discussion explained *why* we can 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 many practically relevant cases, but there are some noteworthy exceptions. Let me mention one example for each condition.

Consider the first condition, 'presence'. 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 this might change. A global nuclear conflict or natural disaster could reduce our planet and the fruits of our civilisation to rubble and dust.[11] If humanity were to survive a catastrophe of this magnitude 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'. 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 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, practice our languages and traditions, and work hard to advance our collective endeavours. If that were to happen and nothing

[11] 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).

of value would be lost in this transition, there would be no conservative reason for extension over replacement either.

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 regularly practised by people, and we already have elaborate systems in place to preserve beautiful works of art: museums and galleries that house them 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 survive alongside us, 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 entirely new cultural artefacts.

To summarise, I have argued 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 elements of our culture. 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 could survive even if our species becomes extinct.[12]

This view has a few further features that are worth mentioning. First, many people share the sentiment that we have a reason to expand humanity diachronically, but not synchronically. The proposed conservative view largely vindicates this conviction. Apart from certain special cases of the sort discussed above, ensuring that another generation exists is necessary and sufficient for preserving the valuable elements of human culture, whereas bringing more people into existence within the same generation isn't.[13]

Second, whether there is a conservative reason to choose extension over replacement seems largely independent of welfare considerations. As we have seen, we could have a conservative reason to choose extension even if replacement would result in a state of affairs that would be more valuable. And presumably we could even have a conservative reason for extending humanity's tenure even if replacement would result in a state of affairs with negative total welfare, if that were necessary and sufficient to safeguard many valuable elements of our culture. These would be merely

[12] I will return to this issue in Sect. 5.

[13] See also Frick (2017) and Lenman (2002).

pro tanto reasons, so they could be outweighed by competing considerations, but they may be a weighty reasons nonetheless.[14]

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-year 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 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, just as in the case of humanity as a whole, 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 non-instrumentally valuable things.

There are two kinds of non-instrumentally valuable things that we should distinguish in this context. On the one hand, there are those which are *good simpliciter*. There 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 these goods. 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 and prosperity could depend on 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.

But there is also another category of non-instrumentally valuable things that we should consider. While discussions of conservatism about value have focused primarily on things which are valuable simpliciter, I argue in other work that we also have conservative moral reasons to preserve *prudential goods*, such as loving relationships and important personal projects, even when a superior replacement is available.[15] On this view, just as there is a pro tanto moral reason to preserve the Golden Gate Bridge even if we could build an even more impressive structure in its place, there can be a pro tanto moral reason to preserve one's marriage or career as a philosopher even when a superior alternative is available.

[14] How weighty? In this, as in many other normative contexts that involve a plurality of reasons or values, a precise answer seems difficult to come by. Perhaps the best we can do to determine what we have overall reason to do in any particular case is to employ the method of reflective equilibrium (cf. Rawls, 1999).

[15] See Masny (unpublished). This view is also endorsed, though not defended, by Nebel (2022).

If this is right, then the conservative account of extension and replacement can potentially apply 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 many people have 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 a person.

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 that warrant preservation are no longer present in a person's life, even if they once were. 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 cannot 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 effort 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.

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. Thus, 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.[16]

[16] Note that conservative reasons are *agent-neutral*: we have them even if we are not personally involved in the relationships or important projects of the person whose life we are in a position to extend. In this

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 Penny will continue to exist for another 5 years. If you instead do Y, Penny will die, but another cow named Rosie will come into existence and live for 5 years. Assume that Penny's extra 5 years and Rosie's 5-year life would not differ in terms of total welfare. Suppose also that doing Y does not amount to killing or failing to rescue Penny. Perhaps you must allocate one unit of a scarce drug, and that drug will either eventually extend Penny's life or else cause Rosie 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 their intuitions about humanity and persons. 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 some of those who accept the practice of 'humane omnivorism', believe otherwise.[17]

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 develop severe depression, and the magnificent cathedral whose renovation she 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 around eight 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.[18]

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 Kolodny (2003, p. 150),

respect, preserving a life which features loving relationships and important personal projects is like preserving the Grand Canyon or the Mona Lisa.

[17] 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 arguments for and against this practice, see McMahan (2008) and Delon (2016).

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

> 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.

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) store food to recover it hours, days, or even weeks later when hungry, which suggests that they can anticipate and plan for future contingencies.[19]

But do these activities count as personal projects of the kind that is worth preserving? According to an influential view, genuine personal projects are similar to loving relationships 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 an animal life 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.

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. This is so even if the replaced animal and the replacement animal would both have excellent lives. 'All' that opponents of humane omnivorism have to do is show that these valuable things are, in fact, present in animal lives.

[19] 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.

Of course, this might not be enough to conclude that humane omnivorism is always impermissible. Conservative reasons are pro tanto reasons, so they must be weighed against moral considerations that might 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.

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 stem, at least partly, from 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 significance of preserving the valuable elements of humanity's culture, such as important traditions, beautiful artworks, and unique languages. To put in different terms, according to the reductive account, 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 unappealing. 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, conservatism should be concerned with preserving the individual's life and humanity (understood as the human species) 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 Frick (2017, p. 359), who asserts the following in his discussion of our 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.

In the same spirit, one might think—as Cohen (2012) and Nebel (2022) appear to—that persons and perhaps animals have the same sort of 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, perhaps non-decisive, reason has to do with theoretical parsimony. As evidenced by the quote from Frick above, proponents of the non-reductive view tend to accept that the valuable elements of our culture are in themselves worth preserving. Likewise, Nebel (2022, p. 188) is at least "inclined to think" that relationships and projects are among the kinds of things we have conservative reasons to preserve. 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 favour the reductive view.

The second reason has to do with extensional adequacy. I believe that the non-reductive view over-generates: there are cases when 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, 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 the capacities 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. In these cases, all that appears to matter to the non-reductivist are brute facts about numerical identity: that the existence of a particular individual or particular species is extended.

I believe that in virtue of the above the non-reductive view betrays the conservative ethos. To make this vivid, take the example of a beautiful painting. Suppose that we can either erase the painting from the canvas and create another beautiful artwork on the same canvas, or else destroy the first painting together with the canvas and create another beautiful artwork on a different canvas. I believe that there is no conservative reason to do the former rather than the latter and I'm confident that proponents of the non-reductive view would share this intuition.[20] Now, I submit that the mere continued existence of a particular individual or species is, in this respect, analogous to the continued existence of a bare canvas. We have no conservative reason to prefer the continued existence of an individual whose life has two discontinuous halves and features two entirely different sets of valuable elements over the consecutive existence of two individuals with corresponding lives.[21] Likewise, we have no conservative reason to prefer the continued existence of the human species whose history has two discontinuous halves and features two entirely different cultures over the consecutive existence of two distinct species with corresponding cultures. This, I think, gives us a decisive reason to reject the non-reductive view.[22]

6 Methods of life extension

In this section, I want to briefly discuss a further issue that concerns certain methods that could be used to extend a person's (or an animal's) life. In recent years, significant progress has been made in our scientific understanding of nutritional, lifestyle, pharmacological, and genetic interventions that can potentially slow down or even reverse ageing.[23] Here, I focus on two of the 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

[20] I set aside cases in which the bare canvas could have certain properties that would render it worth preserving in its own right (say, it's a rare ancient wood panel).

[21] A referee asks: what if what it takes for a person to persist is intertwined with what it takes for the valuable elements of their life to persist? Specifically, perhaps if one were to lose all of their loving relationships and important personal projects, they would cease to be the same person. This is an intriguing view about personal identity, but examining it here would take us too far from the main thread of my discussion. Instead, let me just note that if this view were true, the non-reductive version of the conservative account would have the same implications in the above cases as my preferred reductive version.

[22] A referee wonders whether the reductive view is vulnerable to the 'containers of value' objection traditionally levelled against classical utilitarianism. Let me say three things in response. First, it is important to keep in mind that this objection can be formulated in several substantially different ways (see Yetter Chappell, 2015). Second, the disagreement between the reductive view and the non-reductive view concerns just a narrow slice of normativity. To deny that there is a conservative reason to extend a life devoid of any valuable relationships and projects is not to deny each individual's interests and well-being matter in their own right from the moral point of view. Third, because one's loving relationships and important personal projects cannot be preserved independently of them or redistributed from one individual to another, the reductive view avoids all sorts of implications of classical utilitarianism that many people find unpalatable.

[23] For overviews of recent developments in this field, see Kennedy et al. (2014), Cohen (2015), Newman et al. (2016), and Partridge et al. (2020). In Masny (2023), I argue that we have a justice-based reason to develop healthspan extension technology and make it unconditionally available to everyone.

make revival possible. Uploading, on the other hand, involves scanning a person's brain and uploading this information to a computer, where the person's brain is then simulated.

Concerns about feasibility aside, philosophical discussions of these potential methods of life extensions have focused on two issues: whether these methods would preserve a person's consciousness, and whether they would preserve one's memories and thus their psychological continuity, which many philosophers regard as necessary for personal identity.[24] My discussion of conservative reasons to extend an individual's life puts a spotlight on a third, novel kind of worry. In a nutshell, even if cryonics, uploading, and other similar methods of life extension were consciousness-preserving 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 several 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 might be impossible, or at least more difficult, to maintain these loving relationships. Similar considerations apply to our important personal projects. While these projects 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 especially valuable contents of our lives, even if one's 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.

7 Conclusion

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

According to the conservative account that I have presented, we have a pro tanto 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

[24] For these worries, see Chalmers (2010), Aaronson (2016), and Doyle (2018). For an overview of the literature on personal identity, see Olson (2022).

loving relationships and important personal projects of the kind that warrants further preservation.

Acknowledgements For helpful discussions and comments, I am grateful to Johann Frick, Thomas Kelly, Kacper Kowalczyk, Jake Nebel, Gideon Rosen, and audiences at Princeton University, UC Berkeley, the Moral Theory and the Challenge of Future People Workshop at McGill University, and the Early Career Ethics Network.

Funding 'Open Access funding provided by the MIT Libraries'
The author has no relevant financial or non-financial interests to disclose.

References

Aaronson, S. (2016). The Ghost in the quantum Turing machine. In S. B. Cooper, & A. Hodges (Eds.), *The once and future Turing: Computing the world*. Cambridge University Press.

Arrhenius, G. (2000). An impossibility theorem for welfarist axiology. *Economics and Philosophy*, *16*, 247–266.

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* . 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 P. Thompson, & D. Kaplan (Eds.), *Encyclopedia of food and agricultural ethics*. Springer.

Dickinson, A. (2011). Goal-Directed behavior and future planning in animals. In R. Menzel, & J. Fischer (Eds.), *Animal thinking: Contemporary issues in comparative cognitive*. MIT Press.

Doyle, D. J. (2018). *What does it mean to be human?? Life, death, personhood and the transhumanist movement*. 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. https://doi.org/10.1016/j.cell.2014.10.039

Kolodny, N. (2003). Love as valuing a relationship. *The Philosophical Review*, *112*(2), 135–189.
Lenman, J. (2002). On becoming extinct. *Pacific Philosophical Quarterly*, *83*, 253–269.
Masny, M. (2023). Healthspan extension, completeness of life and justice. *Bioethics*, *37*, 239–245.
Masny, M. (forthcoming). The Shape of History. *Journal of Moral Philosophy*.
Masny, M. (unpublished). Conservatism about Prudential Goods.
McDermott, M. (2019). Harms and objections. *Analysis*, *79*(3), 436–448.
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., & Kennedy, B. K. (2020). The quest to slow ageing through drug discovery. *Nature Reviews Drug Discovery*, *19*(8), 513–532.
Rawls, J. (1999). *A Theory of Justice*. Belknap Press of Harvard University Press, revised edition.
Scheffler, S. (2007). Immigration and the significance of culture. *Philosophy & Public Affairs*, *35*(2), 93–125.
Scheffler, S. (2010). *Equality and tradition: Selected essays*. Oxford University Press.
Scheffler, S. (2018). *Why worry about future generations*. Oxford University Press.
Yetter Chappell, R. (2015). *Value Receptacles Nous*, 49(2): 322–332.