

# A good exit: What to do about the end of our species?

TOBY HANDFIELD

27 July 2017

## Abstract

We know that *Homo sapiens* will not exist forever. Given this, how should our species end? What are the reasons, if any, to delay our extinction? In this paper, I show that the pre-eminent reasons which favour prolonging the existence of the species are partial: they will arise from the particular attachments and projects of the final few generations. While there may also be impartial reasons to prolong the species, these reasons are liable, with time, to reverse their valence: we can be reasonably confident that they will ultimately recommend hastening the demise of the species. Consequently, it is likely that our descendants will eventually face a difficult – possibly tragic – conflict, between partial duties that recommend living on, and an impartial duty to extinguish the species.

This is an author's pre-print of a paper to appear in the *Journal of Moral Philosophy*.

Thanks to Dirk Baltzly, Christian Barry, Steve Clarke, Justin Clarke-Doane, Daniel Cohen, Garrett Cullity, Andy Lamey, Seth Lazar, Barry Loewer, Neil Levy, Chris Meacham, Harry Perlman, Rob Sparrow, Alastair Wilson and numerous audiences for helpful discussions of earlier drafts of this paper. Special thanks to a reviewer for JMP, to Adam Bales, and to Daniel Cohen for very constructive suggestions on recent drafts.

For correspondence: toby.handfield@monash.edu

I think the honorable thing for our species to do is to deny our programming. Stop reproducing, walk hand in hand into extinction – one last midnight, brothers and sisters opting out of a raw deal.

(Rustin Cohle, *True Detective* Episode 1.1 2014)

**1 Introduction** We currently face environmental changes that threaten to radically affect how our species lives. Indeed, it is conceivable that they threaten the very existence of our species. Contemplating these dangers, it is very natural to think – at least for many of us – that extinction of our species would be a very bad thing. So whatever reason we may already have had to address these environmental problems, the risk of extinction appears to give us *further* reason to act. Even if the probability of extinction is very low, the reason it provides to act may still be very stringent, given the distinctively catastrophic nature of human extinction.

An alternative, very different response would be to shrug off extinction as nothing special. This is not to deny that we may have very good reasons to address environmental problems that threaten human wellbeing or other domains of value. Storms associated with climate change will injure people and damage property. Salinization and drought caused by changes in the water cycle could cause food shortages, malnutrition, and famine. It is obvious that we have good reason to try to avoid these outcomes. But that the species may go extinct sooner or later is not itself any further reason for us to take action, on this understanding.

In this paper, I develop an account of the sorts of values that are at stake in our future extinction, which will lead to a somewhat more moderate view than either of these responses. I will argue that although extinction will be very bad for a significant number of people, it need not be utterly catastrophic. In particular, I will argue that the pre-eminent badness of extinction will not be an impersonal evil that makes the world a worse place. Nor will it be a harm that negatively affects the lives of all members of the species. Rather, it will have its greatest impact on those living in the final generations, near the end of our species' history. Because of this, not all humans have similar reasons to avoid a given threat of extinction. The reasons to prolong the existence of the species are partial reasons: they depend upon our particular relationships to the future generations whose existence is at stake, and they will not be shared by other generations.

In addition to identifying these factors that make extinction bad, I will argue that there may be circumstances in which it would be better – on impartial grounds – that the species end sooner rather than later. This leads to the possibility of a tragic conflict. We may have very strong reason to prefer that the species go on, while it would be better, all things considered, to

hasten its demise. Further, the sorts of circumstances that will make extinction desirable are increasingly likely to occur, the longer the species exists. This possibility puts the prospect of human extinction in a rather different light than either of the two responses described above. Rather than assuming that the threat of extinction will always give rise to a demand that present generations make sacrifices to ensure that humans continue exist in future, we must face the possibility that future generations will have moral reasons to make a quite different sacrifice: to ensure that they have no future descendants, even at substantial personal loss.

**2 Death and extinction** To begin thinking about the values threatened by extinction, it is useful to recall why death might be harmful for individuals. Epicurus famously argued that death is nothing to us. He appears to have been impressed by the fact that we are not around to experience our deaths, and this is a reason to think that we cannot be harmed by our own death. This inference should be rejected. Death is a net harm when, had we not died, we *would have enjoyed* sufficient benefits (and not suffered an excessive quantity of harms). In such cases, death deprives us of benefits. That's bad, whether we are there to experience the deprivation or not.<sup>1</sup> For this sort of reason, it is often better to live for a longer period of time than not.

Drawing from the idea that death is usually bad because it deprives us of benefits that we would have enjoyed, had we not died, I suggest that a necessary condition for one particular way of living and dying to be harmful is that it be worse than some alternative ways you might have lived and then died. If a particular death is the culmination of the best possible way of living the rest of your life, then arguably, you have been deprived of nothing by death, and your death is not a harm.

What about for our species? What is the alternative to extinction for a species? We might naively hope that it is simply *non-extinction*: the continued existence of the species, forever. The best evidence, however, suggests that this is not possible. Well established theories in cosmology and thermodynamics entail that the universe will only be habitable by life for a

1. See Bradley (2009: chap. 2) for a similar account of the harm of dying. In particular, Bradley defends a 'Difference-Making Principle': The value of event E, for person S, at world  $w$ , relative to similarity relation R = the intrinsic value of  $w$  for S, minus the intrinsic value for S of the most R-similar world to  $w$  where E does not occur (§2.1). This part of the paper assumes that a similar, though weaker principle, applies to both individuals and species: that for an event E, for person or species S, at world  $w$ , E is bad for S only if there is some reasonably similar world  $w'$  (under a relevant similarity relation R), where E does not occur such that the intrinsic value of  $w'$  for S is greater than the intrinsic value of  $w$  for S.

finite period of time. The nub of the problem is this: eventually, the stars will go out, and the universe will become cold and dark. There will simply not be enough energy available for living organisms to develop. While there are some alternative hypotheses about the ultimate fate of the universe, none of them, so far as I am aware, entail that our species could exist forever (Oppy 2001). So when evaluating the prospect of the species going extinct relatively soon, we must bear in mind that the alternative is not *eternal* human existence; it is simply extinction at a later time.

There are of course very many ways in which the existence of the species could be prolonged, and consequently a variety of different reasons may suggest themselves for why that might be a good thing. In the remainder of this paper, I consider the following types of reasons:

- extending the life of the species might grant some individuals better lives than they would otherwise have had (person-affecting changes);
- extending the life of the species might involve beneficially bringing into existence individuals who would not have otherwise existed (population-affecting changes);
- it might be intrinsically good that the species exist for a longer period of time rather than a shorter; and finally
- it might be instrumentally good, for reasons unrelated to individual welfare, for the species to have a longer existence.

**3 Person-affecting reasons** An initially tempting thought is that the process of extinction is going to be very bad for the individuals who live through it. Imagine something like the worst possible scenario of runaway climate change: massive acidification of the oceans disrupts the entire biosphere; agricultural yields plummet, causing famine; the subsequent food shortages are unevenly distributed, leading to military conflict, the collapse of civilization, the spread of disease, and the eventual extinction of *Homo sapiens*, as well as many other species. Being a member of the last generation of *Homo sapiens*, living through this catastrophe, would be a great hardship. It would be better for those individuals not to live through that process, but to live in more benign times of human history.

Is this a person-affecting reason to prolong the existence of the species? No. We certainly have reason, if it's in our power, to bestow a less awful existence upon the members of this final generation; but this is not a reason to ensure that the species lives for a longer time, rather than a shorter. In other words: we have reason to avoid exposing individuals to famine, disease, and

military conflict – but this is independent of the question whether or not we should prolong the existence of the species.

One way in which the existence of the species could be prolonged is if we prolonged the life of the last human alive. This might be good for this individual: her life might be better, if it were longer. So this is, in some sense, a person-affecting reason to prolong the existence of the species. But it is only ‘accidentally’ a reason to prolong the existence of the species. Were this person not the last member of the species, we would still have just as much reason to prolong her life, even though it would make no difference to the longevity of the species. So we should set aside all such person-affecting considerations for present purposes.

What we require – if there are to be person-affecting reasons to prolong the existence of the species – is that the wellbeing of individuals would be affected in a relatively direct way if the species were to have a longer existence rather than a shorter one. If we hold fixed the events that are relatively local to a person’s life, and vary whether or not *Homo sapiens* goes extinct in 100 years or 1,000 years, would this make any difference to how well that person’s life goes?

In this vein, James Lenman makes the best proposal that I am aware of to explain the dismay that many of us feel about the prospect of imminent human extinction (2002: §4). He claims that it is much like the feeling we would have if, having planned upon having children, we discovered that we could not conceive and raise children of our own. Infertility may frustrate our ability to achieve some of our most cherished hopes. Similarly, for our entire generation, imminent extinction raises the prospect that we will be the last generation, and this suggests that our achievements will be limited in unanticipated ways. Various activities, movements, and enterprises to which we are strongly committed require, for their maximal success, that future generations continue to support and sustain those projects. If the species soon goes extinct, our projects will come to a premature halt. It would be better for us, as members of the final generation, were that not to happen.

Writing more recently, Samuel Scheffler has made similar observations (Scheffler 2012). Scheffler argues that our reaction to discovering that the species faces an imminent demise reveals that the value which many of us find in ordinary activities is contingent upon the existence of others, in the future, valuing similar things. Scheffler believes this is a rational – perhaps even compulsory – feature of the way in which we value things. “We need humanity to have a future for the very idea that things matter to retain a secure place in our conceptual repertoire” (p. 60). If this is correct – and if what matters is *having* a future, not merely believing in a future – then inevitably, for some number of years leading up to the end of our

species, the last generation of humans will suffer a profound loss of value.<sup>2</sup> This is certainly a very serious harm, and there is reason for us to prevent that loss affecting our own lives. I think this is the most powerful sort of reason we have to prolong the existence of the species.

But extinction is inevitable: some generation must suffer the misfortune of having its projects prematurely terminated. Does this mean that, impartially considered, person-affecting reasons cancel out? As Lenman suggests, “from an impersonal standpoint, it makes no very obvious difference [when these evils happen], given that they will happen sometime” (p. 261). Every possible history features an extinction which will typically include this misfortune for the final generation. It is nonetheless consistent with this to say that in every such history, there is a person-affecting reason to defer extinction, because such deferral *would benefit* the final generation. No such person-affecting reasons exist *against* prolonging the existence of the species, because the new generation that would thereby suffer the misfortune of being the final generation would not otherwise have existed. There is thus an asymmetry in our person-affecting reasons that systematically favours prolonging the species.

Person-affecting reasons, however, are notoriously insensitive to the creation of suffering in future generations. So while person-affecting reasons are likely to favour prolonging the existence of the species and these reasons will have an especially strong force for the final generation, we need to consider what reasons might countervail this tendency.<sup>3</sup>

2. How far back will the harm extend? If people abandon the study of philosophy shortly after my death, that seems a much greater threat to the flourishing of my projects than it does to the flourishing of Plato’s similar projects. For convenience, I will talk of ‘the last generation’ as the set of humans who stand to suffer a significant loss of value by a particular extinction. But it should be understood that this set may contain multiple, non-contiguous, biological generations.

In some comments, Scheffler suggests that anything less than an infinite number of descendants furthering our projects will be fatal to value, for *all* preceding generations (Johnston 2014). This seems to be an unnecessary feature of his view, however. Following Kauppinen (2014), I assume that it is important that our projects be pursued for a substantial, but finite period after our deaths, for them to have value. The more extreme view hinted at by Scheffler can nonetheless be accommodated in my framework: it would mean that the ‘last’ generation refers to the *entirety* of humanity to date.

3. I have stressed the badness for *us*, in the final generation, of imminent extinction. This should not be taken to suggest that the person-affecting reason to defer extinction is merely partial. An impartial observer of minimal benevolence would wish to see our projects flourish also. This consideration may have specially strong force for us – they are our projects, after all – but they have impartial weight as well. (I’m indebted to an anonymous reviewer for this point.) That said, I will henceforth refer to these person-affecting reasons as ‘partial’ because it is a convenient way to draw an important contrast with the less ‘counterfactually shifty’ impartial reasons, and because they will almost always, in the sorts of cases in question, have additional force, on partial grounds, for the generation that is in a position to act on them.

**4 Population-affecting reasons** Population ethics is concerned with evaluations of alternative hypothetical populations. For instance, an adequate axiology aims to answer questions such as: is it better, all things considered, that 100 people exist, each with lives of very low welfare; or that 1,000, entirely different people exist, each with lives of very high welfare? In orthodox incarnations, population axiology would rank the second population as better, even though it is better for no individual. Consequently, if it were in our power to bring into existence one or other of these populations, we would have some reason to bring about the second population.

Prima facie, the questions: *how many* people should exist, and for *how long* should our species exist, are distinct. But any decision that is likely to affect the future duration of the species is also a decision affecting which individuals exist in future. A policy decision will initially have impacts on contemporary individuals. But it will also soon begin to have ramifications for who is conceived. After a sufficient period of time has passed – say one hundred years or so – any policy decision will likely bring into existence a set of people, none of whom would have existed had we made any other policy choice. How good is it that said population exists, compared to possible alternative populations? This is precisely the sort of question which population ethics aims to answer. So it is reasonable to expect that population ethics will give us some reason to prefer some longevity-affecting policies rather than others.

Regrettably, the last few decades of work on population ethics have shown that our commonsense intuitions about the value of possible persons are rife with inconsistency. Attempts to capture plausible ideas about the comparative value of populations in formally precise statements lead to contradiction. The only coherent options require giving up one or more otherwise very attractive assumptions.<sup>4</sup> In this essay, I will not attempt to explain in detail the revisionary axiological theories that may be required as a result. Rather, I will employ a rough typology of theories to illustrate the sorts of implications that are likely to arise from population ethics. Population axiologies can be helpfully categorised as *optimistic*, *pessimistic*, or *indifferent*.

*Optimistic theories* tend to treat additional lives as valuable, at least so long as those lives contain net positive welfare. These theories typically have absurd implications regarding the desirability of very large populations. The notorious ‘repugnant conclusion’ is one such example. That is: many otherwise plausible theories entail that, for any moderately large

4. Following Parfit (1984), relevant literature includes: Arrhenius 1999, 2000, 2009; Carlson 1998; Cowen 1996; Huemer 2008.

population, composed entirely of lives of high quality, there is another, much larger possible population, composed entirely of lives of extremely low quality, such that it is *better* that the second population exist. An obvious example of an optimistic theory is total utilitarianism, which equates the value of a population with the aggregate of the utility enjoyed by that population. Other optimistic theories avoid the repugnant conclusion, but at the cost of other absurdities such as the ‘sadistic conclusion.’<sup>5</sup> For instance, average utilitarianism, which states that the value of a population is given by the average of the lifetime wellbeing in the population, entails that it can sometimes be better to add to the population a person with extremely low utility (a life they would prefer never to have lived), rather than add a number of people with modest levels of utility (lives which all concerned would prefer to live).

*Pessimistic theories*, such as negative utilitarianism, regard typical additional lives as having a negative impact on value. Consequently, in their most extreme incarnations, such theories entail that the best possible population is one in which no one ever exists.<sup>6</sup> Less extreme versions such as a maximin theory will still entail that, if we are making decisions under uncertainty, it will almost always be inadvisable to add additional people to the population: doing so cannot raise the minimum that has occurred in the past, but it runs a risk of creating a person with lower wellbeing than has ever been experienced before, thus making the total population worse.

Finally, *indifferent theories* tend to regard any two populations in which different people exist as equal in value or incomparable. For instance, if one adopts a person-affecting view, whereby one population is better than another only if at least one person is better off in one of the populations, then regarding the two populations mentioned at the beginning of this section, we will be driven to the conclusion that neither is better than the other.<sup>7</sup> Such views will also tend to attract the notorious non-identity problem, in which we do no wrong by whimsically bringing about future lives of much lower wellbeing than we could have done.<sup>8</sup>

Having surveyed the terrain in these very simplistic terms, we face an awkward choice with respect to the rest of the argument. Addressing all three possible classes of axiology, in an extended disjunctive syllogism, will quickly become unwieldy. Rather than do that, I propose to set aside pessimistic and indifferent views (though personally I think both these perspec-

5. Arrhenius (2000); Ng (1989).

6. Benatar (2006); Fehige (1998).

7. Handfield (2011); Roberts (1998).

8. Parfit (1984).



tives deserve significantly more credence than they are typically given).<sup>9</sup> I make this choice because, given that my overarching argument is that our eventual extinction is not likely to be a catastrophe, it will be of greatest dialectical interest if I can support this conclusion against a background assumption that most readily favours increased longevity.

Even within the optimistic camp, there is a great diversity of theoretical positions we could adopt. For concreteness, I assume that the following two claims are common to all credible optimistic theories.

**Better-with** For any population of individuals, all of whom are above a sufficiently high threshold of wellbeing, it is better that this population exist than not.

**Better-without** For any population of individuals, all of whom are below a sufficiently low threshold of wellbeing, it is better that this population not exist.

In the previous section, I indicated that person-affecting reasons can be insensitive to suffering that occurs in future generations, because typically, the only alternative to such suffering – for that generation – is that it not exist at all. But what if the suffering is so great that the entire generation would actually prefer never to have existed? In that case, is it reasonable to think that we have *harmed* the generation? I am deliberately silent on this issue. *Better-without* entails that we can make things impartially worse by bringing future people into existence if their lives are sufficiently wretched, but it is consistent with this claim that none of those individuals are actually harmed. Similarly, it is consistent to suppose that they are harmed, and that the harm is the basis for the world being impartially worsened by their ex-

9. Some recent work has attempted to use decision theoretic reasoning over this matter of normative uncertainty, to determine whether some theories give rise to reasons which, even if we give them very low credence, will swamp all other possible reasons. Arguments of this sort would suggest that we should ignore the advice of indifferent theories, for instance, because they will say that it does not matter what we do: they generate reasons that are always swamped (Ross 2006). In unpublished work, Greaves and Ord (2015) apply this approach to all axiologies that assign invariant cardinal values to possible populations, including total and average utilitarianism. The result favours a view like total utilitarianism rather than average utilitarianism, because as populations become very large, the reasons generated by a total view swamp any effect of the average. But this approach relies upon the highly implausible assumption that there is a meaningful quantity in terms of which inter-theoretical comparisons of value can be made – precisely the issue that makes normative uncertainty so troubling. And even if that assumption can be granted, Greaves and Ord’s argument leads to the conclusion that the theory to guide our decisions will be a version of critical level utilitarianism. That is, a view on which lives above a critical threshold of utility make a positive contribution to total value, and those below make a negative contribution. The exact critical level depends on how our credences are distributed over the possible axiologies they consider. Negative utilitarianism is arguably an extreme case of critical level utilitarianism: one where the threshold is infinite, and hence their conclusion is precisely a hybrid of pessimistic and optimistic theories.

istence. I suggest we can remain agnostic as to this issue. Provided we grant the claim about impartial worsening, we have sufficient reason to predict that there can be tragic conflicts of the sort identified above.

4.1 *Tragic conflict between generations* These assumptions vindicate the suggestion made in the introduction. We can face a decision in which we will benefit ourselves by averting the harm of being a generation whose projects are prematurely abandoned, but we will thereby create new generations, some of whom will suffer this fate instead. Moreover, creating those new generations could make things worse, all things considered, if the lives created are of sufficiently low wellbeing.

To understand the range of issues such choices might generate, it will be helpful to develop a toy model of the way in which value is constituted by projects being carried out over generations, and the way in which this value is enjoyed by each generation. Suppose that the world's population can be divided into a finite number of non-overlapping generations  $P = \{1, 2, \dots, n\}$ .<sup>10</sup> Each generation  $i \in P$  enjoys a level of wellbeing  $w_i(n)$  which is an aggregate of 'private' goods  $p_i$  that are produced and enjoyed within that generation, and 'shared' goods  $s_i(n)$  that are obtained from *later* generations participating in the ongoing enterprises and projects of the present generation. The quantity of shared goods, unlike private goods, depends on the number of future generations and their efforts to support or frustrate the ongoing projects of the present generation.

$$w_i(n) = p_i + s_i(n)$$

This makes clear the way in which a generation can be benefited by prolonging the species: assuming that additional future generations will create some shared goods that will

10. The enjoyment of shared goods presumably depends upon the existence of meaningful relationships between generations. For instance, if the cultural achievements we hope that our descendants will attain were instead spontaneously produced by an unrelated species of aliens, in a remote part of the universe, it would make no contribution to our wellbeing. It is not essential that future cultural achievements be produced by members of our biological species for us to benefit from them, but rather that there be some cultural continuity between our achievements and those in the future. Because I am supposing that there will be no species other than *Homo sapiens* that achieves cultural continuity with present generations, I can, without mischief, conflate the notions of cultural and biological extinction. But if we were forced to distinguish, the bulk of the arguments in this essay suggest that we should be concerned primarily about cultural extinction. So we assume that all the generations in the model are culturally contiguous. The model could be made more sophisticated by introduced a measure of relatedness between generations, such that the contribution to shared goods varies with the weight of the connection.

contribute to the value of the lives in the present generation (i.e. the local derivative of  $s_i$  is positive with respect to  $n$ ), then a generation can, by reproducing, benefit itself by increasing  $n$  and in turn increasing the value of the shared goods term.<sup>11</sup>

I assume that both private goods and shared goods can have negative valence. Other things being equal, it would be better if such ills never existed.<sup>12</sup> To give some plausible examples: if we assume that private goods include hedonic experiences, we can readily imagine future generations having lives that are filled with excruciating pains and physical hardships, such that it would have been better for those experiences never to have occurred. Regarding shared goods, it may be helpful to imagine that the achievement of just institutions in future is a shared good – future generations continuing to uphold justice will benefit us. Conversely, if our descendants perpetrate vicious injustices, the ignominy of this might be thought to detract from the value of our own lives. Another possible example of shared value is aesthetic achievement. If future generations destroy and vandalise the greatest artworks of previous generations, and in their stead create hideous works of kitsch, we might think that this diminishes not only the value of their lives, but ours also.

Granted these preliminaries, the following possibilities arise:

- If the private goods that would be enjoyed by the next generation are positive and large ( $p_{i+1} \gg 0$ ), but the shared goods they would produce for the present generation are negative ( $s_i(n) > s_i(n + 1)$ ), then we may have partial reason to prefer that the next generation not exist, while from an impartial perspective it would be better if they did exist.
- If the private goods enjoyed by the next generation are negative and large ( $p_{i+1} \ll 0$ ), but the shared goods they would produce are positive ( $s_i(n + 1) > s_i(n)$ ), we may have partial reason to bring the next generation into existence, while from an impartial perspective it would be better if they did not exist.

11. This model can also be understood to represent the way in which we sometimes appear to take sunk costs into account in our actions, acting in order to ‘redeem’ or ‘vindicate’ sacrifices made by others in the past (Kelly 2004). If an action undertaken now is intrinsically more valuable because it redeems a past sacrifice, this is equivalent to an action generating a shared good that benefits earlier generations as well as present.
12. To remain neutral on the controversies of population ethics, it should not be assumed that we can calculate the impartial value of creating the  $i$ th generation by simply considering the value  $w_i$ . However, we can assume – following Better-without – that for *sufficiently* low levels of  $w_i$ , assuming everyone in that generation had sufficiently low wellbeing, it would be impartially better if generation  $i$  did not exist. For convenience, stipulate that  $w_i$  is negative only when this condition is satisfied.

The second case is the sort of tragic conflict we identified above. Creating future generations may be permissible, given the very significant partial reasons that we possess to ensure that our valuable enterprises continue, but if the conditions under which those generations exist are very adverse, it may be all things considered worse that they existed.

(If it is legitimate to call the creation of such wretched generations a *harm* to those generations, then the sense of tragedy becomes all the sharper: to save ourselves, we must harm innocent others.<sup>13</sup> If, as many believe, there are moral constraints against harming innocent others in order to benefit oneself, this suggests it may be impermissible to prolong the species in such a case (Otsuka 1994). Extinction would be compulsory.)

This sort of conflict is not merely possible: it is almost certain to arise, if our species lives long enough. Note first, the issue arises only when *all* future generations are most probably faced with unrelentingly grim conditions. It is not enough to suppose that the next ten or twenty generations will face adversity, because, given the very modest claims we have adopted from population ethics thus far, we cannot rule out the possibility that adversity in the immediate future will, in the long run, be outweighed by the flourishing of yet later generations. But because we know a good deal about the destiny of our ecosystem on cosmic timescales, we can be quite confident that there will eventually come a time when the conditions for human existence will be unrelentingly grim. The most obvious difficulty is the finite life of our Sun. In the later stages of its life, around five billion years from now, the Sun will swell to become a red giant, literally engulfing the Earth. For millions of years prior to that, the Earth will be too hot to be habitable anyway. Even if the species escapes the solar system prior to death of the Sun, a similar fate ultimately awaits our descendants in any similar solar system that they colonise. If this is not grim enough, on even longer timescales, the increase of entropy poses an even harder limit to the pleasantly habitable regions of the universe.<sup>14</sup>

13. Writing on the moral problems posed by climate change, Stephen Gardiner (2006) identifies a similar structure. If earlier generations do not take sufficient steps to limit climate harms, then subsequent generations may need, in order to avert harms to themselves, to take steps which exacerbate climate-related harms for subsequent generations. In this way, a chain of threat-relations could be triggered between generations: each vulnerable generation is innocent of any wrongdoing, but must generate a threat to a subsequent generation in order to protect its own vital interests. Thanks to Greg O'Hair for illuminating conversation of related issues.
14. This answers an objection raised by a reviewer: the impartial reasons I am offering that support an earlier extinction are in some sense merely 'accidental'. They might be more appropriately characterised as reasons to make the lives of future generations more pleasant, rather than reasons to extinguish the species. This is true, provided that we have a range of options available that are likely to make the lives of future generations more pleasant. But eventually, for the reasons indicated, it will be inescapable that prolonging the species will make

(More immediate dangers, such as climate change and nuclear warfare, could also be so severe in their effects that a similar issue might arise. But even on the worst such scenarios, there is likely to be a non-trivial chance that the species will outlast the adversity involved, and we must therefore – on current evidence – keep an open mind about the possibilities for later generations to eventually flourish. All of this is conditional, of course, on some variety of *optimistic* population ethics. Pessimistic views entail that the future is unrelentingly grim already.)

It is also interesting to contemplate whether any credible scenarios could give rise to the first sort of case, where partial reasons favour ending the species, while impartial reasons favour prolonging it. Perhaps we might foresee a future in which our cherished cultural institutions will be destroyed and replaced by practices that we regard as hateful.<sup>15</sup> Our dread for this possibility could be so great that we prefer that the species not continue. But we might simultaneously recognise that the population which could exist in future would enjoy a substantial quantity of private goods. These goods alone might warrant saying that it would be better, all things considered, if that generation were brought into existence. I take this sort of case to be rather less plausible than the other, because it strikes me that it is improbable for there to be such sharp discontinuities in cultural values between generations. A generation whose practices are so hateful to us may not have much claim to being our actual cultural *descendants* at all, and hence if all future generations are like this, our culture may have already been effectively extinguished.

So, on the assumption of a broadly optimistic approach to the value of possible populations, there is indeed some reason to prolong the existence of the species: namely, we may thereby increase the aggregate of valuable lives lived, and we may benefit ourselves by averting the imminent abandonment of our projects. But whether or not extinction is a catastrophe, on this view, depends very much on how good those future lives would be. Plausible approaches to population axiology are not relentlessly optimistic: they admit the possibility that future lives could be so wretched that they are better not lived. (Indeed: the key claim – *Better-without* – is common to both optimistic and pessimistic views.) Further, given what we know about the ultimate fate of the cosmos, if our species manages to remain alive long enough, there will come a time when all future lives are certain to be wretched. In a circum-

---

things impartially worse. In such circumstances, hastening extinction is the only available way to avert wretched lives for future generations.

15. Johnston 2014.

stance like this, partial reasons to prolong the species may still obtain, while impartial reasons favour imminent extinction.

My first aim has been achieved: I have shown that one of the most significant reasons for avoiding extinction is partial, and that impartial reasons will eventually support hastening the extinction of our species. But the argument is not complete. I have so far focused solely on reasons related to welfare and human flourishing. To fully assess the plausibility of the idea that extinction need not be a catastrophic event, we need to consider other potential ways that species longevity might be valuable.

**5 Species-level reasons to prefer longevity** In this section, I consider the possibility that, independently of any facts about individual wellbeing, it might be intrinsically better that the species have a longer existence.

Longevity is a matter of the temporal distribution of individual lives in the species. Theories of population axiology pay heed to how much aggregate welfare there is, and how it is distributed across persons, but it is blind to *when and where* the lives occur. Is there reason to think that the temporal or spatial distribution of lives in a population matters?

Similar issues have been raised in discussing the nature of individual wellbeing. It is very plausible to think that what makes an individual life go well is not merely the aggregate of episodic benefits, less any episodic ills, but that the temporal distribution of benefits and ills makes a further contribution to wellbeing. David Velleman (1991) makes an argument of this sort, premised on the idea that human lives have a sort of narrative structure which entails that some distributions of benefit are preferable over others, even if the aggregate quantities of benefit are equivalent. For instance: it is better to have a life that goes out on a high, having started at a low, rather than to have the opposite trajectory.

This sort of reasoning can be brought to bear directly on the value of longevity. Consider Figure 1, which illustrates two different sorts of life: one is a 'live-fast, die-young' life that contains many goods packed into a shorter lifespan, versus another life that contains a similar quantity of goods stretched out over a longer period of time. Perhaps, if we focused purely on worldly achievements as the sorts of goods in question, the live-fast, die-young life could be imagined as a rather idealized version of the life of a prodigious high-achiever, like Mozart, who died at age 35, having composed approximately 600 works. And let us suppose that the second life contains a similarly impressive raft of achievements, but was stretched out over a more normal lifespan of 70-odd years. This could be likened to the life of Haydn, who wrote approximately 470 compositions, and lived till 77 years of age.

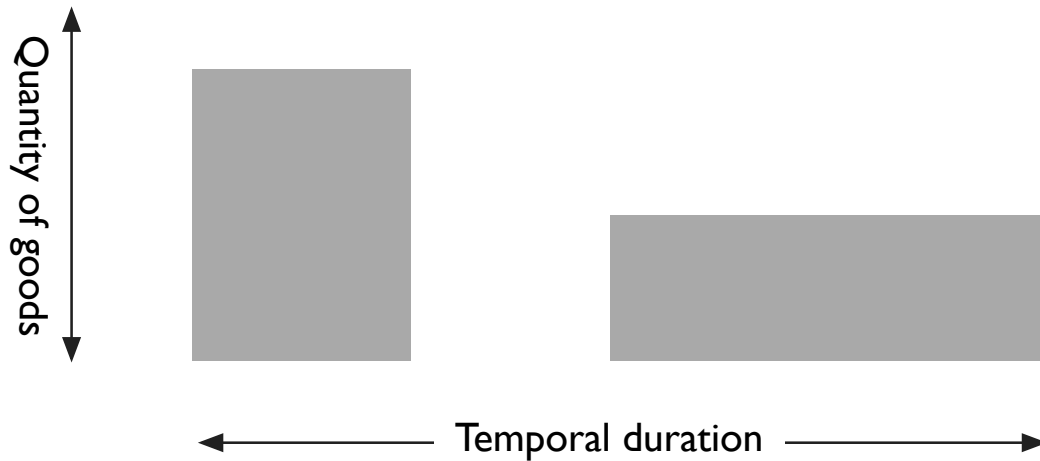


Figure 1: Two lives: one a ‘live-fast, die-young’ type of life, and another, more normally distributed life.

Putting aside any idiosyncratic preferences one might have for the music of Mozart or Haydn, there is some plausibility in the thought that, of these lives, other things being equal, the longer one is better. That would suggest that a longer life is not only valuable because a longer life allows us to experience or obtain more good things than a shorter life. It suggests, further, that long life might be somewhat good *in itself*. This is at least a prima facie plausible claim. This idea might be especially attractive to those who endorse some sort of Aristotelian conception of the good life: there might be a certain span of life that is simply the proper span for a species of our sort. Lives shorter than that span are, other things being equal, worse.<sup>16</sup>

Taking inspiration from this idea, is there any reason to suppose that it is better that the ‘shape’ of a population be longer in the temporal dimension rather than shorter? I am sceptical that this is so. As James Lenman writes:

We might have taken it to have such a structure if we had some large philosophical vision of human history as making sense in terms of some readily discernible goal which it might be tragic not to attain. I take it very few of us today are

16. John Broome employs a different thought experiment to identify the value of longevity (Broome 2004: §7.2): two short lives versus one long one. Although I share Broome’s intuition that the single, longer life is preferable, it is undesirable to rely on a thought experiment which requires a comparison between scenarios in which different people will live – thereby forcing us to contemplate questions about goodness for *populations*, rather than merely goodness for an individual.

gripped by such a vision. If human beings go on for countless millennia, today will seem to have been the childhood of our species. If we disappear tomorrow, today will seem (to some imaginary observing aliens) to have been its old age. If we reject grand philosophical pictures that endow human history with some essential pattern, all that can be meant by metaphorical talk of our species' childhood is those times that are *relatively* early in its career whenever they may turn out to be. The individual human tragedy of dying young has no obvious analogue in the career of our species as a whole. (Lenman 2002: p. 259)

To make this a bit more concrete, compare two possible populations of humans which can be thought of as alternative histories for the entire species (Figure 2). One is a fast-breeding population that rapidly comes to achieve a very high population density. The average population at any time is approximately 8 million (around the population of New York City). This population lasts for only 100 years. The second species is a much more long-lived population which maintains a steady population of approximately 160,000 (around the population of Swindon, England; Geelong, Victoria; or Tempe, Arizona) and endures for 5,000 years.<sup>17</sup> Of course there are some benefits to living in a less dense population: less stress, less noise, less traffic, cleaner air. And there are some benefits to living in a large city-type environment: more varied employment opportunities, better restaurants, better public transport, better live entertainment, and a generally larger range of goods and services to purchase. But we are to imagine, for the purposes of this example, that there is a correspondence between the quality of the lives lived in the New York population over the 100 year period and lives lived in the Swindon population over the 5,000 year period: there is a one-to-one mapping between lives in the two populations, such that for each mapped pair, neither life is better than the other. So the benefits and drawbacks of big and small city living are counterbalanced in the two populations. The only relevant difference between the lives lived in the two scenarios is that, in the New York case, the lives are bunched closely together in time, and in the Swindon case, they are spread out over a much longer temporal interval.

In making claims about intrinsic value, it is very difficult not to beg the question against one's opponents. But for what it is worth: other things being equal, I see no reason to prefer one of these species-histories over the other. Neither species has an obvious claim to be

17. And in neither case, remember, are we to suppose that the population comes to an end in an especially horrible way: that would be to focus on a particular *process* of extinction rather than the value of longevity *per se*.





Figure 2: Two populations – a dense and short-lived species and a sparse and long-lived species.

flourishing more than the other. So I conclude that there is no intrinsic value in longevity for a species.

**6 Instrumental reasons to prefer species-longevity** Finally, I consider other non-welfare-related goods that might be furthered by prolonging the existence of the species.

*6.1 The end of progress* It has sometimes been put to me that, although there is no narrative structure to human history, there is something inherently good about the upward trajectory we have observed in living standards, or in technological advancement, or in the state of collective knowledge. So in some sense, progress is itself good, and this is thought to provide a reason to avoid the end of the species.

In claiming that progress is good, the advocate of this argument could mean a number of different things. We can set aside the thought that progress with regard to a valuable enterprise is instrumentally good because it enables individuals to enjoy more episodic goods derived from that enterprise. That hypothesis is subsumed by the discussion of welfarist reasons above.

A second possible interpretation is that it is intrinsically better that progress go on for a longer period rather than a shorter period. On any reading that I can conceive, this proposal implausibly fetishizes the gradient of achievement over time, rather than valuing achievement itself. It is hard to credit that it would be worth forgoing any substantial quantity of achievement in exchange for increasing the period of time over which the level of achievement was increasing. Consider the alternatives: humans could discover 100 chemical elements, but dis-

tribute that achievement over 1,000 years. Or we could discover 105 elements, but over 900 years. It is difficult to credit that the former is better, on account of the longer duration of progress.

A third – and much more plausible – interpretation is that it is better that, at some point in time, we attain a higher rather than a lower level of achievement in the relevant field of endeavour. The ‘high point’ of human civilization is in some way relevant to the value of the species’ existence overall. What the lover of progress intrinsically values then, is the highest possible achievement. A greater period of existence is instrumental towards this end because it provides additional opportunity to reach a higher level of achievement than we otherwise might.

While I do not find this perfectionist proposal particularly plausible, I cannot offer an outright refutation. It is important to note, however, that this idea gives only qualified support to increased *longevity* for the species. Suppose we reached a plateau of achievement in the relevant field of endeavour, and because of resource scarcity, we will never exceed that level of achievement. This argument would now give no reason to prefer a longer existence for the species. Indeed, one of the few advocates of perfectionism of this sort, Thomas Hurka, implies that our achievements may be devalued if they are followed by additional lesser achievements.

Many of us think Muhammad Ali’s boxing career ... would have been better without those last fights against Larry Holmes and Trevor Berbick. This is not because we think Ali’s performances against Holmes and Berbick were by some objective standard bad; we know that, for many other boxers, to do as well as Ali did against these fighters would have marked the pinnacle of their careers. It is rather because we think Ali’s performances were so much worse than the performances he produced in his prime that it was bad *for him* to produce them. The Holmes and Berbick fights were mere additions to Ali’s boxing career, yet many of us think his career would have been better without them.

(Hurka 1983: 505–6)

Hurka’s comments appear to imply that, if a plateau is reached in the realm of human achievement, it would actually be preferable that the species come to an end earlier, rather than undergo a decline in its material or cultural standing.

It might be insisted that the relevant domains of achievement are *not perfectible*, so there will always be room for further improvement. It could also be insisted that we will never be in a position to *know* that we have reached a peak of achievement: we will at worst have diminishing confidence in the possibility of further achievements. The force of these objections

will of course depend on the nature of the achievements which we take to be relevant, and I prefer to avoid such controversies here. But for the same sorts of reasons that we can be sure the species will eventually go extinct, I believe there will be definite *physical* limits on how much our species can achieve, in any domain. The universe will, eventually, become hostile to further achievements, no matter how good our intentions may be, and we may thereby come to have very high credence that we have reached our limits. This is entirely compatible with their being no logical limit to achievement in our domains of endeavour.<sup>18</sup> So even though perfectionism may provide reasons — for much of human existence — to prefer that our species exist for a longer period, these reasons are at some point guaranteed to diminish to zero, and may – on Hurka’s type of approach – favour hastening extinction.

6.2 *The decay of cultural artefacts* Human culture has produced numerous astonishingly beautiful, complex, intricate, and spectacular artefacts: symphonies, paintings, pop songs, novels, sculptures, buildings, gardens, poems, and so on. Many of these are *good*. Most obviously, they are good because the enjoyment of these artefacts constitute episodic goods in people’s lives. But let us put aside those considerations because they have already been addressed in the discussion of individual welfare above. Could there be other values, not reducible to the enjoyment of episodic goods, attached to the existence of these artefacts?

As above, it may be tempting to argue that:

- P1. The existence of these artefacts is good.
- P2. The extinction of our species will lead to the decay and destruction of these artefacts.
- C. So extinction is bad.

But as before, it is important to frame the argument properly to identify the relevant alternative. In all possible histories, the gradual increase in entropy will lead to the destruction of all cultural artefacts. So a better argument might be:

- P1. The existence of cultural artefacts for a longer time is intrinsically better than the existence of cultural artefacts for a shorter time.
- P2. The extinction of our species sooner rather than later will lead to the decay and destruction of cultural artefacts sooner rather than later.
- C. So extinction of our species sooner rather than later is bad.

18. It remains an open possibility that the universe is infinite in size, and although this would allow that an infinite amount of information can be stored, there are limits on how much of that information could be a causal product of our culture. Thanks to Barry Loewer for advice here.

Now that the issue is properly framed, the question is the plausibility of the first premiss. One reason why it might appear plausible is because we assume that, if cultural artefacts exist for longer, there will be more episodes in which people can enjoy them. But obviously that is not relevant in this context, where humans are assumed to be extinct before the artefacts are destroyed. Moreover, this does not support the intrinsic betterness of artefact-longevity.

So if the first premiss is true, it is good, independently of whether any humans exist to enjoy them, that cultural artefacts exist a longer time rather than a shorter. This is rather implausible.<sup>19</sup> And even if it were conceded, the value of artefact-longevity is surely not very significant. Suppose we had a choice between two scenarios:

**Scenario A.** The human species will go extinct in 200 years time, due to a massive meteorite strike. All cultural artefacts will be vapourised at the same time.

**Scenario B.** The human species will go extinct in the very same scenario as A, except for two differences: one person will suffer an additional, intense but short-lived pain, such as being kicked hard in the shins; and the contents of the Uffizi Gallery in Florence will be preserved from destruction by the meteorite, and will thus last for an additional 200 years.

I take it that the preservation of the contents of the Uffizi Gallery for an additional 200 years would, other things being equal, be a good thing, according to P1. But would it be sufficiently good that it would outweigh the badness of one person suffering a painful kick in the shins? I find it extremely difficult to accept that it could be. So even if there is intrinsic value in the preservation of cultural artefacts, it is of trivial weight compared to the avoidance of harms to individual persons. Consequently, the prospect of preserving cultural artefacts provides us with very little reason to delay our extinction, relative to considerations of harm and benefit to persons.

6.3 *Conservatism and valuing particular things* Dismay at imminent extinction might be grounded in the sort of valuing articulated by G. A. Cohen as a *conservative* attitude, which recommends conserving particular valuable *things*, rather than merely conserving the value that is in things (Cohen 2011). Even if we could replace a valuable particular with something of equivalent or greater value, the conservative has some reason – admittedly a defeasible reason – to resist this, and to retain the less valuable particular that is already possessed.

19. The same thought is voiced by Lenman (2002: 255).

This conservative attitude could be extended towards our cultural practices, languages, artefacts, or biological substrates. These things are bearers of value, and – notwithstanding that they are going to cease to exist eventually; notwithstanding that they might, in some scenarios, be replaced by more valuable alternatives – conservatism regarding such objects *would* provide a robust reason to regret imminent human extinction, because it would entail the loss of these things. Such an attitude could provide additional reasons, beyond any so far considered in this essay; or it could constitute an alternative diagnosis of the attitudes discussed by Lenman and Scheffler. The Lenman and Scheffler-type reasons are reasons to prefer that projects are pursued for a longer rather than a shorter time, because the value of our lives, or of our projects, is enhanced by such pursuit. But a conservative take on this preference for continued pursuit of projects could be simply that these are *our* projects (or they are already *existing* projects – Cohen considers both ideas as variants on conservatism), which are valuable, and we regret any abandonment of them, even if they are replaced by projects that we anticipate will ultimately be even more valuable.

Cohen's conservatism is not the same as valuing longevity: it does not entail a positive attitude towards long-lived states of existence. It is more akin to a deontological constraint, requiring the prioritization of efforts to prevent destruction of existing valuable things – this constraint obtains regardless of how long a thing has existed to date, and of how long it will continue to exist in future. Unlike a more traditional deontological constraint, it does not have especially agent-centred content: it does not require that *I* avoid destroying valuable things.<sup>20</sup> But it clearly has a non-teleological character: it directly prescribes particular sorts of behaviour towards valuable things, rather than aiming to promote or maximize value.

I won't attempt here to defend or critique this particular type of valuing. Rather, I rely on the claim that, whatever additional reason may be found on these grounds for delaying extinction, it is surely – as Cohen admits – defeasible. Consequently the impartial reasons that I anticipate arising to favour the demise of the species will in some circumstances outweigh the conservative reasons to delay it. While the conservative reason should not become significantly more weighty as time progresses, we can anticipate that in the distant future there will be a significant increase in the magnitude of the reasons to hasten extinction. At such times, conservative reasons will be swamped.<sup>21</sup>

20. See Cohen, pp. 218–9.

21. A similar reply can be given to those who reject my assertion that species longevity is not intrinsically valuable. I'm indebted to Garrett Cullity for some of these points.

6.4 *Loss of biodiversity and other species conservationist reasons* There is some appeal to the notion that, while life exists, other things being equal, it is better that there be a rich, bio-diverse array of species in existence at any given time, rather than a narrower, impoverished range. The extinction of a species at an earlier time rather than a later time will typically entail that biodiversity is lower, for a time, than it would have been. So this looks like a reason for thinking that the imminent extinction of a species is bad.

This consideration is quite unlike the earlier proposals, which tried to identify something bad about extinction for a species. On this proposal, the extinction of a species at an earlier time is an instrumental evil, because of its impact on biodiversity, which is not a good enjoyed by a species, but is a good that inheres in an entire ecosystem. Even if correct, however, it is limited in what it says about the evils of extinction. First: it is silent about the extinction of all life. If all species were to go extinct simultaneously, this would not constitute a loss of biodiversity in the relevant sense. This provides relatively little comfort then. Second: if we were to adopt a stronger claim, that it is better that life exist in a biodiverse state for as long as possible, that would amount to asserting that long-lasting ecosystems or biospheres are intrinsically valuable. Like the claim that there is intrinsic value in longevity for a single species, I find this implausible. The personal dismay we might feel at witnessing the imminent extinction of a species does not give very much support to the thought that it is intrinsically better for biodiversity to be long-lasting. It is too readily explained by the person-affecting hypothesis that *our* lives are impoverished if we live in a less biodiverse environment than we might have done.<sup>22</sup>

One further argument can be offered against this stronger biodiversity view: it appears to imply the absurd idea that it would be desirable to resurrect the dinosaurs and have them living now, on grounds of enhancing biodiversity. The ideal state of the world would supposedly be one in which we kept the richest possible menagerie in existence for as long as possible.

Finally, as Lenman notes, appeal to biodiversity in support of the idea that extinction is an evil will work only for *some* extinctions. Species that have an adverse impact on biodiversity will perhaps best be extinguished. And the prime candidate for a species that is having an adverse impact on biodiversity is of course, *Homo sapiens* (Lenman 2002: 257). Biodiversity will surely fare much better if we are gone as soon as possible.

22. I am inclined to offer a similar diagnosis of the intuition that the death of a language is intrinsically regrettable (Levy 2001).

There are a number of other arguments made for species conservation in general, and I cannot hope to do justice to them all. I note, however, that the dialectical space is significantly limited once we focus on *human* extinction. Stereotypically, arguments regarding species conservation are made to persuade a human audience that there is reason not to destroy another distinct species. Many such arguments have little or no force, when the question is whether to take steps to delay the extinction of our own species.

For instance, Ian Smith argues that the virtue of *humility* should motivate our behaviours that impact on other species. Humility involves ‘recognizing that not everything of importance is in what happens to oneself and one’s close friends and family’ (Smith 2010: 179). Humility, then, could play an important role in bolstering arguments for *hastening* our extinction. A concern for humility may lead us to recognise that we are placing too much emphasis on our essentially partial concern to avoid being among the final generations of humans, as discussed in section 3. Relatedly, we may come to downplay our concern to preserve the human species, once we see the terrible cost it inflicts on numerous other species on the planet.

To take a second type of argument for species preservation, Bryan Norton claims that natural phenomena have *transformative* value. Our experiences of nature can transform our preferences and values, and do so in ways that make those preferences more rationally defensible or worthy. The extinction of a species, then, deprives us of a potential source of this valuable transformative experience (Norton 1987). While this argument may have considerable appeal for why we – organisms who can experience a transformation of values – should conserve other species, it does not generalise to the preservation of our own species. Human beings are not sources of transformative value for any other species, except perhaps in cases of domesticated animals. But it is difficult to believe that the preferences we instil in our pets for human affection or for play with other humans are more rationally defensible or worthy than the preferences those animals would have had without human contact.

**7 Conclusion** Would human extinction be a catastrophe? This question is too stark to be of much use. Better question: in what ways would extinction be bad? I have identified three central ways: (i) it may harm the final generations by reducing the quantity of shared goods they enjoy; (ii) extinction may deprive the world of additional lives that were worth living; and (iii) extinction may limit the degree of human achievement. But it is important to note the corollaries of these answers. In virtue of the population-related concerns cited in (ii), extinction may be a blessing because it prevents a generation of wretched lives coming into

existence. The perfectionist reasons alluded to in (iii) may entail that it is better that we go out of existence before human degradation and mediocrity becomes yet more commonplace, sullyng our past achievements. And finally, because the person-affecting considerations cited in (i) are to some extent partial, and acutely felt by each generation when it exists, we can predict that if a time comes when we are better off extinct, there may be some motivated blindness to this fact. It may require great moral courage to identify that it is time to go.

#### REFERENCES

- Arrhenius, G. (1999). An impossibility theorem in population axiology with weak ordering assumptions. In Sliwinski, R., editor, *Philosophical Crumbs*. Department of Philosophy, Uppsala University, Uppsala.
- Arrhenius, G. (2000). An impossibility theorem for welfarist axiologies. *Economics and Philosophy*, 16(02):247–266.
- Arrhenius, G. (2009). One more axiological impossibility theorem. In Johansson, L.-G., Österberg, J., and Sliwinski, R., editors, *Logic, Ethics, and All That Jazz*. Department of Philosophy, Uppsala University, Uppsala.
- Benatar, D. (2006). *Better Never to Have Been: The harm of coming into existence*. Oxford University Press, New York.
- Bradley, B. (2009). *Well Being and Death*. Clarendon Press, Oxford.
- Broome, J. (2004). *Weighing Lives*. Oxford University Press, Oxford.
- Carlson, E. (1998). Mere addition and two trilemmas of population ethics. *Economics and Philosophy*, 14(02):283–306.
- Cohen, G. A. (2011). Rescuing conservatism. In Wallace, R. J., Kumar, R., and Freeman, S., editors, *Reasons and Recognition*. Oxford University Press, Oxford.
- Cowen, T. (1996). What do we learn from the repugnant conclusion? *Ethics*, 106:754–75.
- Fehige, C. (1998). A Pareto principle for possible people. In Fehige, C. and Wessels, U., editors, *Preferences*, pages 508–43. de Gruyter, Berlin.
- Gardiner, S. M. (2006). A Perfect Moral Storm: Climate change, intergenerational ethics and the problem of moral corruption. *Environmental Values*, 15:397–413.
- Greaves, H. and Ord, T. (2015). Population ethics and normative uncertainty. Paper presented at Ethics and Decision Theory Workshop, ANU.
- Handfield, T. (2011). Absent desires. *Utilitas*, 23:402–27.
- Huemer, M. (2008). In defence of repugnance. *Mind*, 117:899–933.



- Hurka, T. (1983). Value and population size. *Ethics*, 93:496–507.
- Johnston, M. (2014). Is life a Ponzi scheme? *Boston Review*.
- Kauppinen, A. (2014). Flourishing and finitude. *Journal of Ethics and Social Philosophy*.
- Kelly, T. (2004). Sunk costs, rationality, and acting for the sake of the past. *Nous*, 38(1):60–85.
- Lenman, J. (2002). On becoming extinct. *Pacific Philosophical Quarterly*, 83(3):253–269.
- Levy, N. (2001). Why regret language death? *Public Affairs Quarterly*, 15(4):373–84.
- Ng, Y.-K. (1989). What should we do about future generations? *Economics and Philosophy*, 5:235–53.
- Norton, B. G. (1987). *Why preserve natural variety?* Princeton University Press, Princeton.
- Oppy, G. (2001). Physical eschatology. *Philo*, 4(2):148–68.
- Otsuka, M. (1994). Killing the innocent in self-defense. *Philosophy and Public Affairs*, 23:74–94.
- Parfit, D. (1984). *Reasons and Persons*. Clarendon Press, Oxford.
- Roberts, M. A. (1998). *Child Versus Childmaker: Future Persons and Present Duties in Ethics and the Law*. Rowman & Littlefield.
- Ross, J. (2006). Rejecting ethical deflationism. *Ethics*, 116(4):742–768.
- Scheffler, S. (2012). *Death and the Afterlife*. Oxford University Press, Oxford.
- Smith, I. A. (2010). The Role of Humility and Intrinsic Goods in Preserving Endangered Species. *Environmental Ethics*, 32(2):165–182.
- Velleman, J. D. (1991). Well-being and time. *Pacific Philosophical Quarterly*, 72:48–77. Reprinted in Velleman 2000: 56–84.
- Velleman, J. D. (2000). *The Possibility of Practical Reason*. Clarendon Press, Oxford.