Utilitas (2024), **36**, 211–229
doi:10.1017/S0953820824000098

CAMBRIDGE UNIVERSITY PRESS

ARTICLES

Totalism, Animals, and the Repugnant Conclusion

Gary David O'Brien

Department of Philosophy and Hong Kong Catastrophic Risk Centre, Lingnan University, Hong Kong
Email: garyobrien@ln.edu.hk

Abstract

Totalism states that one population is better than another iff it has higher total welfare. One counterintuitive consequence is the Repugnant Conclusion (RC). Totalism also entails that a very large population of animals with lives barely worth living is better than a smaller population of happier humans. Furthermore, the strategies that have been used to avoid the troubling normative implications of the RC do not work in the animal case, so we may have reason to bring about such a population. I introduce the notion of 'Efficiency of Welfare Production' – that animals of different species vary in the efficiency with which they convert resources into welfare. If we want to maximize total welfare, without any speciesist bias, we should identify which species is most efficient and try to maximize the population of that species. This has counterintuitive implications whether we accept hedonism or a more sophisticated theory of welfare.

Keywords: Totalism; Population Ethics; Repugnant Conclusion; Animal Ethics

Section 1

1.1 Introduction

Since Parfit's (1984) *Reasons and Persons*, a huge literature on population ethics has emerged. A key aim of work in this area has been to avoid the Repugnant Conclusion (RC). Some population ethicists have argued that we should endorse the RC (Huemer, 2008; Tännsjö, 2002). Others have agreed that avoiding it should not be a criterion for an acceptable population axiology (Zuber et al., 2021). A notable omission from the literature is the lack of attention paid to nonhuman animals. This may simply be an oversight, as most ethicists are more concerned with humans than with animals. Alternatively, population ethicists may have assumed that considering non-human animals will not make a significant difference to their theories.

In this paper I will explore how the inclusion of animals bears on population ethics, paying particular attention to the RC. While in principle the RC has always been species-neutral, most discussions of it in the literature have focused exclusively on human populations. This has led population ethicists to overlook certain implications that only arise when we remember to include nonhuman animals in our considerations.

We can think of the standard RC in traditional population ethics as if it included only possible *human* populations in its scope.[1] Likewise, the Z world typically referred to a world in which there was a very large *human* population composed of individuals with lives that are barely worth living. To flesh out the implications of considering animals in our population ethical thinking I will use the term animal-Z to refer to a world in which there is a very large population of relatively simple nonhuman animals each of whom has a life that is barely worth living.

My use of this terminology should not be construed as suggesting that animals are not, or have not, been included in the scope of population ethics or the RC. Nor am I claiming that my paper is original in 'discovering' that Totalism can be applied to animals. Rather, I use this term to make it clear that there are differences between typical humans and non-human animals which are relevant to our population ethical thinking, especially when we are thinking about populations we could actually bring about, rather than about populations that are purely hypothetical. The first difference is that small animals are much more resource efficient than humans, and, consequently, we could create and sustain a much larger population of animals than of humans given the same amount of resources. The second difference relates to the typical quality of life that is available to humans and nonhumans, in particular to the possibility that humans, but not animals, can enjoy the so-called 'higher pleasures' or 'best things in life'. I discuss the implications of these differences in 3.1 and 3.2.

Totalism is the view that one population is better than another iff it has higher total welfare (Greaves, 2017). In this paper I will understand Totalism to also include anti-speciesism – that is, to reject the claim that the value of a unit of welfare varies depending on the species of the individual in whom that welfare is realized. I will also understand Totalism as endorsing a welfarist axiology – that is that welfare, and only welfare, is valuable.[2] If we reject speciesism, then the best population according to Totalism may be one composed of nonhuman animals. That is, animal-Z may be more valuable than an A population of humans. If this conclusion is more counterintuitive than accepting that a Z population of humans is better than the A population then this puts pressure on Totalism. This is an *Axiological Problem* of combining Totalism with anti-speciesism. I will argue that this combination also raises a *Normative Problem*, that we may have reasons to try to bring about something like an animal-Z population, and this may be more challenging than the Axiological Problem. I will argue that there are differences between humans and animals which make avoidance of the Normative Problem much more difficult in the animal case than in the human case (2.2). If I am right, then this poses a separate, perhaps more serious, challenge to Totalism.

If we cannot avoid the Normative Problem in the animal case, then, it seems that we have reason to bring about something like an animal-Z world, if this is the best way to maximize total welfare.[3] I argue that we could in fact best maximize total welfare by bringing about something like an animal-Z world (3.1). I do so by introducing the notion of *Efficiency of Welfare Production* (EWP). This is the idea that animals of different species vary in the degree of efficiency with which they convert resources into

[1]Though see Williamson (2021) and Sebo (2023) for recent exceptions.

[2]This is not the only way of understanding Totalism, but it is a natural one. If the reader disagrees, they are welcome to replace 'Totalism' with 'Totalism*' to indicate the view I am exploring.

[3]My argument assumes the possibility of interspecies welfare comparisons. This is, of course, a big assumption. See Browning (2023) on the complexity and imperfect nature of such comparisons.

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

welfare. I argue that some species of animal have higher EWPs than humans do, and so a population composed of these more efficient species would have higher total welfare than a human population sustained by the same amount of resources. Assuming simple hedonism, the best possible population in terms of welfare is a vast one composed insofar as possible of members of whatever species turns out to have the highest EWP. I suggest mice as a possible candidate for the most efficient species. Of course, determining which species is most efficient at converting resources into welfare is a complex question requiring much more empirical and philosophical work, and it is likely that some other species will turn out to be more efficient. My choice of mice should be understood as illustrative of the general idea of EWP, not as the strong claim that mice are in fact the most efficient species. In any case, this population of small, relatively simple animals would have very low average welfare and it would lack humans and their distinctive pleasures. I argue that it would also lack species diversity and wild animals generally.

Finally, in section 4 I consider the objection that we could avoid this conclusion by adopting a richer theory of welfare that gives special weight to those activities distinctive of humanity, such as listening to excellent music, engaging in intellectual pursuits and so on. A theory of welfare that gives extra weight to those things Mill called the higher pleasures, and Parfit referred to as the best things in life could identify humans as the species with the highest EWP, and so would support the creation of a population composed insofar as possible of human beings.[4] Though this population is appealing in some ways, it is also seriously lacking in species diversity and wild animal life generally. If Totalism implies that one of these two populations is the best, and yet neither is intuitively appealing to us, this may give us reason to doubt Totalism. If, on the contrary, one finds these populations acceptable, then this paper can be seen not as an argument against Totalism, but as an exploration of some of its underdiscussed implications.

1.2 The repugnant conclusion and the normative problem

The RC is seen as a challenge as it is very difficult to accept that the existence of a Z world in which vast numbers of people live lives that are barely worth living could possibly be better than the A world in which 10 billion people live excellent lives. Parfit's original understanding of Z was that it is a drab world in which individuals have few pains which are barely outweighed by the simple pleasures of 'muzak and potatoes', and from which all the greater accomplishments and ecstasies of excellent human lives are absent (1986, p. 148). On this understanding of Z, it is indeed very counterintuitive that it could be better than A, no matter how many people Z contains. Meeting the challenge has proven extraordinarily difficult, however.

Some philosophers have endorsed the RC (Huemer, 2008; Tännsjö, 2002) and others have agreed that avoiding it should not be a necessary condition for an acceptable population axiology (Zuber et al., 2021). If one endorses the RC and accepts that Z is more valuable than A then a new problem arises, however. The problem is that if we acknowledge that Z is more valuable than A, and we accept that we have at least a prima facie moral reason to bring about more valuable states of affairs when it is in our power to do

[4]Sebo (2023, p. 261) raises the possibility of beings with a capacity for welfare far higher than ours – something like Nozick's Utility Monster. If such beings existed, then the best population in terms of welfare might be a smaller population of these creatures, rather than a larger population of humans. Since such creatures are purely hypothetical, I will not consider this possibility further.

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

so, then it seems that we have a prima facie reason to try to bring about a Z world, or at least something approximating it should that be practically possible. Even if we are willing to endorse the RC in our axiological thinking, many of us would still find it difficult to believe that we have any moral reason to try to bring about such a state of affairs. If we initially accept the RC in our purely axiological thinking, then go on to reject the implication that it gives us any reason to bring about Z like populations, we might see this as a reductio of the axiological claim. We can call this *the Normative Problem* of the RC to distinguish it from the axiological problem.

Why, however, should we think that the Normative Problem is a separate problem beyond the axiological one? After all, if one genuinely believes that the Z world is more valuable than the A world, then it shouldn't be difficult to accept that we have reasons to try to bring about such a world should that be possible. There are several reasons why the Normative Problem may be important. First, a dialectical reason. Since all theories in population ethics have counterintuitive axiological implications, in deciding between competing theories we may look to features other than their axiological implications to make our decision. If Totalism has both counterintuitive axiological implications and counterintuitive normative ones, while some other theories have only counterintuitive axiological implications, this may be a reason to disfavor Totalism. Note that I am not claiming that this is the case. It may turn out that all theories in population ethics have both normative and axiological implications that are counterintuitive, and in that case, we would be back to a dialectical deadlock. I only make the case here that, once we remember to include animals, Totalism does in fact have implausible normative implications. Second, if we accept the methodology of reflective equilibrium, then the fact that a given theory has counterintuitive results at more than one level (here, axiological and normative) counts against that theory, if there are other theories that have counterintuitive implications at only one level. This is because all our intuitions are taken to be evidential, not only the more foundational (axiological) ones. Finally, at least for some people, normative intuitions about what we have reason to do are stronger than highly abstract intuitions about what we ought to value in hypothetical (and perhaps practically impossible) cases. If we give more weight to normative intuitions about what we ought to do in possible cases than we do to our axiological intuitions about impossible ones, then a theory that has implausible normative implications is, all else being equal, weaker than one that does not.

1.3 Avoiding the normative problem

Though they don't use the term, both Tännsjö and Huemer seem to recognize the Normative Problem, and gesture at strategies to make it less troubling. Huemer explicitly asks whether, if we accept the RC as he urges us to, we ought to "aim at a drab future like world Z, where each of our descendants occupies a single, cramped room, and there is just enough gruel to keep them from hunger?" (p. 928). He argues that, given any plausible view about the actual effects of population growth on total and average wellbeing, Totalism will not recommend trying to maximize the number of people with lives barely worth living. Rather the best way to maximize total welfare will be to reach the optimum population, which is likely to be a point at which there is a large population at a still relatively high, though not maximum, average level of wellbeing (p. 929). For example, if we doubled the world's population, this might result in the average wellbeing level being much lower than 50% of current levels, thereby reducing total wellbeing too. Totalism therefore would not imply that we ought to maximize the population as that would reduce total welfare relative to the optimum population size.

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

Furthermore, Huemer points out that it is not clear that increasing the population will reduce average wellbeing at all. While larger populations mean less space, and a lower share of resources for each person, they also have positive effects on average welfare. More people could mean more innovation, more art, a greater pool of talent to draw from and so on. If these positive effects of large populations on average welfare outweigh the negative effects, then the RC is rendered practically toothless. The apparent trade-off between total and average welfare would turn out to be an illusion, and though the axiological problem posed by the RC would remain, in practice it would be side-stepped by the positive welfare effects of a large population. Having made this response to the Normative Problem Huemer concludes that Totalism "does not enjoin us, in reality, to pursue the world of cramped apartments and daily gruel" and suggests that critics of Totalism "will therefore look upon the principle with less revulsion than has hitherto been customary" (p. 930). This final remark suggests that Huemer believes that at least part of the reason that people find endorsing the RC repugnant is that it seems to have such unpalatable normative implications. We can label Huemer's response to the Normative Problem the *Unpredictable Effects Argument*, as it questions the assumption that we know what the effects of population increase on welfare will be.

Tännsjö also wishes to correct the mistaken belief that accepting the RC means that one should also endorse all increases in the world's population, as some increases may reduce total welfare. He also reminds us that the RC describes a merely hypothetical situation, and that the population in Z is, in Parfit's own words, an *imaginable* one (Tännsjö, p. 342). This suggests a further way of blocking the Normative Problem which we can label the *Impossible Populations Argument*. The Z population is merely imaginable because of its magnitude. Of course, there is no fact of the matter about how many people are in the Z population, but we can make reasonable estimates. If the A population has ten billion people as Parfit stipulates, and the people in the Z population have lifetime wellbeing levels 100 times lower than those in the A population, then for the total wellbeing in Z to be greater than that in A the Z population would have to be more than 100 times larger. That is more than one trillion people, and it is not at all clear that the planet can sustain that many human beings.[5] The implication is that, since the massive Z population is entirely hypothetical, we need not worry about the normative implications of endorsing the RC.

Another possible response to the Normative Problem is what I will call *the Control Problem*. If we endorsed the RC and thought that we had reason to try to bring about a very large population of human beings with lives barely worth living, we would have to exert very significant control over human reproduction indefinitely to create and maintain it. This would of course be almost impossible to implement. Furthermore, if states coerced people to have more children this would likely have severe negative effects on the wellbeing of the population, possibly sufficient to counter the positive welfare effects of creating extra people. Finally, even if it were practically possible to successfully implement such extensive control over people's reproductive decisions, many people recognize deontological constraints against such infringements of people's autonomy. If this is correct, then it is neither practically possible, nor morally permissible, to take the measures required to create a Z like population. Even if we accept that Z is more

[5]The UNEP (2012) has produced a survey of 65 estimates of the planet's maximum carrying capacity. Only one of the estimates is as high as 1 trillion people, and most estimates are between 8 and 16 billion people.

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

valuable than A, if it isn't practically possible or morally permissible to move us toward Z then we can not have a moral reason to do so. The RC then would be an awkward, counterintuitive fact about what we ought to value, but it would not have implications for how we ought to act. Endorsing the RC need not mean endorsing 'repugnant actions'.

Section 2

2.1 The RC and animals

Though his original formulation of the RC concerned only human beings, Parfit recognized that the same logic that led us from excellent lives to muzak and potatoes could be extended across the species barrier, resulting in Z worlds composed of nonhuman animals (2016, p. 118). A few others have written on the possibility of extending the RC to include nonhuman animals. Tännsjö recognizes that the RC really ought to be stated in terms of sentient beings rather than people and claims that it might be speciesist to reject a world with higher total welfare on the basis that the population is composed of nonhumans (2002, pp. 339, 352). Williamson (2021) uses the animal case as an argument against Critical Level axiologies. Critical Level views are those according to which lives worth living can only contribute to the value of a population if they surpass the critical level. Williamson argues that such views cannot accommodate animals, as they run into insurmountable problems whether they make the threshold universal or they relativize it to species. Sebo (2023) argues that expanding our moral circle to include animals raises questions about priority setting and what sort of population we have reason to aim at. Assuming that small animals typically have lower average welfare levels than larger ones, and that large populations of small animals might have higher total welfare than a small population of large animals, Sebo argues that, at least in the long run, we may have reasons to try to bring about a population that has more, but smaller animals rather than one composed mostly of larger animals. He calls this 'the rebugnant conclusion' as on his view we should also include insects and perhaps even microbes in our moral circle (p. 11).

If we reject speciesism, accept the possibility of interspecies cardinal comparisons of welfare, and endorse a Totalist welfarist axiology it is easy to see how we can get to the animal-Z population. If we start with a population A consisting of 1 million happy humans, we can see that population B consisting of 2 million chimpanzees who are each more than half as happy as the humans in A is better, as total welfare is higher. From B we can move to C, a population of 4 million dogs, each of whom is more than half as happy as each of the chimps in B.[6] This operation can be repeated indefinitely, until we reach a vast population of simple animals with a very low capacity for happiness. Perhaps this population would consist of trillions of barely sentient oysters, each of whom has a life that is just barely worth living. If the happiness in each oyster life in Z is a million times less than the happiness in each human life in A, then, as long as there are more than a trillion oysters in Z, the total happiness in Z will exceed that in A.

[6]The species and numbers chosen here are merely for illustrative purposes – one need not believe that chimps have half the lifetime wellbeing levels of humans, or dogs half that of chimps to understand the move from the A population of happy humans to the Z population of simpler animals. It would also be possible to have a more gradual transition from A to Z by including our hominid ancestors, e.g. having a B population of *Homo erectus* and so on.

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

As Williamson (p. 410) points out, for those who found the original human-focused RC repugnant, the animal-inclusive version is likely to be an unwelcome result. If it is repugnant to conclude that a population of humans living excellent lives is less good than a much larger population of humans with lives that are barely worth living, it must be at least as repugnant to conclude that a sufficiently large population of simple animals would also be better. It may be that most people have the intuition that animal-Z would be even worse than human-Z, yet Totalism implies that it is as good or better, and this may reduce the plausibility of Totalism. I am not convinced, however, that animal-Z must be more troubling to Totalists than human-Z.

When we are considering populations that are merely logically possible, it is not clear that animal-Z need be any worse than human-Z. In thought experiments we can make whatever stipulations we like about the welfare levels of the inhabitants of Z, whether human or animal. For example, we can stipulate that the average level of welfare is arbitrarily low, that each life lasts only a few moments as in Short-Lived Z, that there are no higher pleasures, but only 'muzak and potatoes' as in Drab Z etc. In thought experiments the differences between a human-Z population and an animal one can be stripped away, until the only salient difference between them is species. If we reject the claim that the value of a unit of welfare depends on the species of the individual in which it is instantiated, however, then there is no reason to prefer a human-Z population over an animal one. Furthermore, accepting that animal-Z is better than a human A population does not automatically entail that we have moral reason to bring about something like an animal-Z population. One might think that the responses to the Normative Problem of the RC can be applied equally well to the animal-inclusive case, thus blocking the move from a counterintuitive result in axiology to an even more counterintuitive account of our reasons for action. In the next section I will argue that this supposition is false. There are relevant differences between humans and animals which mean that the strategies for responding to the Normative Problem in the human case do not apply in the animal case.

2.2 Animal-Z and the normative problem

We saw in 1.3 some ways in which those who endorse the RC in their axiology can avoid the Normative Problem. I will argue in this section that these strategies for avoiding the Normative Problem posed by the RC don't work when we remember to include animals. Consider first Huemer's Unpredictable Effects Argument. His claim was that endorsing the RC axiologically doesn't necessarily entail that we should maximize the human population, rather we should seek the optimum population, which may be one with a relatively high average welfare. When we include animals in our considerations however, things are different. First, as I will argue in greater detail in section 3, since small animals have much lower resource requirements than humans, we could support vastly higher populations of animals, and this would probably increase total welfare. Secondly, it is plausible that most simple animals have lower lifetime welfare levels than most humans. Most human beings have very good lives compared to animals like mice.[7] Humans have a wide, diverse, and complex array of pleasures, deep emotional lives, and sophisticated intellectual interests, while mice, presumably, don't.

[7]In describing mice as 'relatively simple' throughout this paper, I do not intend to imply that mice are objectively simple, merely that they are simple compared to humans. Mouse minds, brains, and behavior are still objectively complex.

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

Furthermore, the average human lives for around 80 years, while the average mouse lives for about two years. In fact, to think that mice and humans have approximately the same level of lifetime welfare despite this difference in life expectancy, we would need to think that the welfare level of a mouse at any given moment is many times higher than that of a human, and this does not seem plausible. The lifetime wellbeing of a mouse, even a very happy one, may be hundreds of times lower than that of a human being. Bringing about a Z population of mice then would certainly lower average wellbeing – the trade-off between average and total may be illusory in the human case, but in the pan-species case the trade-off is real.[8]

Second, consider Tännsjö's observation about the impossibly large population in the human-Z world. It may be correct that a human-Z population is too large to sustain on Earth and must remain purely imaginary. This objection does not seem to hold in the animal case, however, since we can sustain very large populations of small nonhuman animals. Consider the number of animals humans farm globally each year as an illustration. According to estimates by Fish count (2019a, 2019b), in 2017 between 250 and 600 billion crustaceans were farmed, as well as 51–167 billion fish. We also slaughter an estimated 50 billion chickens each year, and several billion sheep, pigs, and cows (Thornton, 2019). These numbers illustrate the practical possibility of creating and maintaining very large populations of nonhuman animals. If we raised smaller animals it seems likely that we could support even larger populations, easily rivaling those of the hypothetical human-Z population.

Of course, the lives of animals raised on factory farms are not good ones, and, may not be worth living – raising a vast population of animals in such conditions would not maximize total welfare and may even result in a population of negative welfare value. Indeed, part of the explanation for why we can raise such large numbers of animals is precisely because they have low welfare – factory farms are designed to optimize the production of animal products, and this goal comes at the cost of animal welfare. The important question for my purposes is not whether we can sustain huge nonhuman populations, but whether we can do so while also ensuring that the animals have net positive welfare. This is a difficult empirical question, but it does not seem impossible to sustain very large populations of small, relatively simple animals with net positive welfare. First, keeping mice as our example, there is a huge disparity in size, and resource requirements between chickens and mice, and so all else being equal, we should be able to raise many more mice, even at a higher welfare level, than chickens.[9]

Secondly, the requirement that we can raise large numbers of animals at high population densities will affect our choice of most efficient animal. If mice cannot thrive at the high population densities required to maximize total welfare, then we may choose another animal that can. Perhaps there are equivalently small, relatively simple animals who flourish in dense populations. Alternatively, if we are sufficiently confident that simpler animals such as insects or oysters are capable of positive welfare, we might choose to breed vast numbers of them, instead. Insects of many kinds live naturally in dense populations, and so it seems unlikely that this is detrimental to their welfare.

[8]Huemer's argument might suggest that the optimal mouse population is a smaller one of happy mice, rather than a much larger one of mice with lives barely worth living. However, even a population of happy mice is likely to have much lower average wellbeing than any realistically possible population of humans. Small, longer-lived animals might have relatively high lifetime wellbeing though.

[9]An average laying hen is 100–200 times heavier than a mouse and consumes 30–50 times more food each day.

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

Whatever animal we choose, the upshot is that if we deliberately aim to optimize for total welfare, rather than for efficiency of food production, it is plausible that we could sustain truly massive animal populations with net positive welfare.

Finally, one might ask whether we would be *willing* to sustain such large numbers of happy animals even if we were able to. I do not claim that we have any self-interested reason to do so. I merely point out that, if we are Totalists, then we need to think seriously about how we can best maximize welfare, and it might turn out that sustaining large populations of nonhuman animals is the best means to that end. Nonetheless, even if supporting a large population of small animals is possible in principle, it could be impossible to implement for other reasons. This could be because most people are not committed Totalists, and creating and sustaining a huge population of small animals would require global agreement and co-operation, which is unlikely to be forthcoming. It might then be the case that the impossible populations argument blocks the normative problem in the animal case as well as the human case, though the impossibility here is not physical, but social impossibility. However, in the human case the normative problem is blocked by the unpredictable effects argument, the control problem, and the impossible populations argument, while in the animal case only the latter argument is effective. Still, those who are committed Totalists could increase welfare by maintaining large numbers of small animals themselves even if this falls short of their ideal population (see 3.2). Furthermore, if implementation of the Z population of small animals is impossible, this might give Totalists reason to prefer the 'human-world' I describe in section 4.

Finally, consider the Control Problem. While it would be very difficult or impossible to control human reproductive behavior sufficiently to bring about something resembling a Z population, it seems relatively easy to control the reproductive behavior of some animals. Questions of the permissibility of such control seem less salient in the animal case too. While it would be monstrous to interfere significantly with the reproductive decisions of humans, this seems less problematic in the case of simpler animals like mice. Humans make autonomous decisions about their reproductive lives based on many considerations about what they value and how they want to live their lives. Mice do not seem to make reproductive decisions in the same way that humans do – rather their reproductive behavior seems instinctive, and so there would be less of an issue with undermining the autonomy of mice by controlling their reproductive behaviors. However, we might think that animals like mice, though far simpler than humans, do have some interest in self-determination, and that we might wrong them by controlling their reproductive lives in this way.[10] If this wrong is only slight, then it may be outweighed by the greater welfare we generate by controlling animal reproduction. It is also likely that even simpler animals, like insects or oysters, have little or no interest in self-determination, and, for animals like these, controlled reproduction need not be technically nor morally troubling.

Sebo (2023, pp. 9–12) argues that, if we can maximize total utility by creating many small animals with low average levels of welfare, then we may have reason to do so, at least in theory. He recognizes that this is a counterintuitive conclusion, which many might want to avoid. In my terms, Sebo recognizes the Normative Problem, though he does not think it is necessarily a conclusion that we should try to avoid (p. 12). Nevertheless, he offers two considerations which might allow us to block the practical normative implications of extending the RC to animals, even if we still accept the

[10]See Healey and Pepper (2021).

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

axiological claim that the animal-Z world might be better than our own world. We can consider these as solutions to the Normative Problem that are specific to the animal case.

The first solution concerns the tractability of creating a large population of very small animals and ensuring that they have good lives. Sebo argues that this may be much more difficult than creating smaller populations of happier large animals. Though he thinks it would be tractable to create more insects and other small animals, it might be much more difficult to ensure that they have net positive experiences, especially if they continue to reproduce through r-selection.[11] This suggests that Sebo has in mind the creation of large numbers of *wild* animals; r-selection is only a welfare problem for wild animals, as they quickly reproduce beyond the carrying capacity of their environments. If small r-selecting animals were bred in captivity, then it would be relatively easy to provide them with sufficient resources, and to control their breeding so that they don't produce more offspring than can survive. Sebo might also think it might be difficult to ensure that insects have good experiences, since they are so different from us. This is an epistemological problem – as we move further away from humans on the tree of life, our confidence in our ability to measure and promote welfare in other animals might reasonably waver. Both problems can be solved if we raise small mammals in captivity. This would solve the welfare problems caused by r-selection, as we could effectively control the reproduction of captive animals to ensure their population growth does not outstrip available resources. It would also mitigate the epistemological problem as we can be much more confident in our assessment of the welfare of animals that are evolutionarily closer to us than we can be with insects. I will investigate this possibility in greater detail in section 3.

Second, Sebo argues that when we consider the long-term effects of our actions it may no longer be optimal to create large populations of small animals in the short term. Rather, we might best promote welfare in the long term by focusing our efforts on improving the lives of human beings as, in the long run, humanity will be in a better position to maximize welfare than we are today. This suggests that, even if prioritizing the welfare of smaller animals is the best way to maximize welfare in the present, we are permitted to prioritize humans for now as this will have the best consequences in the long term. As Sebo acknowledges, however, this is not really a solution to the Normative Problem. Considering the optimal strategy for maximizing welfare in the long term "might save us from a surprising neartermist implication, but it would do so only by replacing it with a surprising longtermist implication" (p. 13). Furthermore, it is not clear that even in the short term the best way to maximize total wellbeing in the future will be to focus on human beings. Some ways of helping humans today will increase our future ability to help small animals, but others will not. For example, saving the lives of current day humans, or lifting them out of poverty might instead lower total wellbeing, if those humans will consume animal products from factory farms (Plant, 2022). Furthermore, as Beckstead (2013) has argued, not all human beings will be able to have positive effects on the long-term trajectory of humanity, and, in the case of such people, the Totalist might prefer to spend resources on creating large numbers of happy animals in the present.

[11]The reproductive strategy followed by most animals on Earth. It means having many offspring but investing little energy in each one. It poses a welfare problem, as most of these animals die painfully shortly after birth. See Horta (2010) and Johannsen (2017).

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

The upshot of these arguments is that while the human-focused RC may be merely a theoretical problem in axiology, the animal case also poses a Normative Problem. Totalists could safely ignore the RC when deciding how to act since it is impossible to actualize something like the human-Z population. Blocking the natural move from the axiological to the normative makes endorsing the RC, and thus accepting Totalism, less unappealing. But since it is possible to bring about something like an animal-Z population, and since this also arises naturally from Totalism, this may count against the plausibility of Totalism. In the next section we will investigate what a practically implementable animal-Z world might look like.

Section 3

3.1 Population engineering 1: mouse world

Let us assume the position of an ambitious would-be 'population engineer' who wishes to maximize total welfare in a non-speciesist way. Since she has limited resources, she needs to think about what kind of population will most efficiently maximize welfare. Most importantly, she will need to determine whether there are any salient differences between species that will make some better generators of welfare than others. I suggest that there are such differences, and which species are most efficient for welfare generation will depend on how we understand welfare. If we think that all sentient beings have the capacity for wellbeing, and that there are no sharp discontinuities between the kinds of wellbeing available to humans and those available to simpler animals, then we might think that the best possible population is one composed of small, relatively simple animals such as mice or oysters.[12] The reason for this is straightforward. Small simple animals generally require very few resources and little space, while large and complex animals like humans require a large amount of both. Since we inhabit a planet with finite resources, if we are to maximize the welfare value of the population then we have good reason to fill the world with small relatively simple animals rather than large complex ones, so long as the total amount of happiness in the 'small animal world' would be greater than that in the 'large animal world'.

In this section I will assume this conception of welfare is correct. Given her goal of creating a population with the highest welfare value, her rejection of speciesism, and the limited resources available to her, how should our population engineer proceed? An obvious method is to identify which species of animal (including humans) generates the greatest balance of positive wellbeing per unit of resources, and then attempt to move us towards a large population that is composed of individuals of that species to the extent possible. Human lives are generally high in welfare, but the cost of maintaining human lives is also high. It may turn out that there are simpler animals who can 'convert' resources into welfare more efficiently than humans do. Let's call this conversion rate 'Efficiency of Welfare Production' (EWP).

Since we have limited resources, we must consider EWP when thinking about what sort of population to bring about. EWP is likely to vary not only among different ways of producing welfare, and between different individuals of the same species, but also between species. This is because different animals may have comparable capacities for welfare, while some require much fewer resources than others. In particular, given plausible assumptions about the levels of welfare than nonhuman animals are

[12]This conception of welfare is endorsed by Bentham (1789), Sidgwick (1907), McTaggart (1927), and de Lazari-Radek and Singer (2014).

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

capable of, human beings are inefficient welfare producers. Humans are large, complex animals. Compared to many smaller and simpler animals, they require a huge amount of resources to survive and thrive. Smaller animals may attain similar levels of welfare as humans do, while requiring far less food, space, and other resources.

In the rest of this section, I will compare the EWP of human beings with that of mice, and I will argue that, given plausible assumptions about the wellbeing experienced by mice, we have reason to try to bring about a large population of happy mice instead of a much smaller population of happier humans. I have chosen mice for several reasons. Firstly, because their physical size means that they require much less food, space, and other resources than humans do. Secondly because they are mammals. This isn't because there is any greater value in mammal happiness than happiness in non-mammals. Rather it is for epistemological reasons. While there are certainly many simpler animals than mice, once we leave the mammalian order, doubts about their sentience, and the kinds and degrees of wellbeing available to them, become much greater. No one reasonably doubts mouse sentience, and it seems that they enjoy many kinds of wellbeing which humans share and understand. While it is certainly possible that we could produce more happiness by creating a very large population of animals simpler than mice this is less certain given our current understanding of sentience and wellbeing, and the physical structures required for their realization. Of course, if one is more confident about the welfare capacities of even simpler animals, and of our ability to promote their welfare, then an even stronger argument can be made using those simpler animals in place of mice. Sebo (2023) argues that, at least in theory, we should focus our efforts on even simpler animals such as insects, nematodes, and even microbes. This is because even if our credence that they are sentient is low, there are so many of them that the expected value of trying to increase their welfare is very high. I will continue to focus on mice as they are sufficiently efficient welfare converters to make my argument, but do not require the additional complexity of discounting for uncertainty, or gambling on high but uncertain expected values.

In order to compare the EWPs of humans and mice we need to know how much welfare each is capable of realizing, and what quantity of resources it takes for them to do so. This requires having a theory of welfare, and being able to measure and compare the extent to which humans and mice are capable of enjoying it. Let's assume simple Benthamite hedonism as our theory of welfare. We then need to ask how a mouse's capacity for pleasurable experiences compares to that of a human. Clearly it is not possible to measure intensity of pleasurable experiences directly. However, we can get a clearer picture by thinking about the degree of pleasure mice and humans experience from engaging in the same kind of activity. Consider the pleasure a human gets from eating a meal. Let us assume that the meal generates 1 unit of pleasure. Now imagine how much pleasure a mouse gets from eating a meal. Certainly, it is plausible that mice take pleasure in eating, for much the same evolutionary reasons that humans do. But how much pleasure does a mouse take in eating? Approximately the same amount? Ten times less than a human does? One hundred times less? It seems implausible that the pleasure the mouse gains from her meal could be any less intense than that. So, all else being equal, if we had the choice of giving a meal to a human or to a mouse, we ought to give it to the human. Not on speciesist grounds, but simply because the amount of pleasure generated is likely more than the mouse would experience, and so we can best maximize welfare by feeding the human rather than the mouse. However, all is not equal. On average a human being eats 1.8 kg of food per day. A mouse eats 3 grams. Even if the human being's pleasure in her meal is 100 times greater than the

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

pleasure the mouse takes in hers, a human eats 600 times more than a mouse. If we want to maximize pleasure, without any speciesist bias, then we have reason to invest our limited resources in producing food for mice instead of food for humans, as this will be at least six times more efficient, even when ignoring the fact that the financial cost of producing food palatable for humans is much higher than producing the same amount of food palatable to a mouse.[13]

This suggests that, at least when it comes to the pleasure of eating, mice have a higher EWP than humans do, unless we make implausibly low estimates for the amount of pleasure mice get from eating. Of course, it isn't just the pleasure of eating that we are trying to maximize, but overall welfare. This too is plausibly maximized by directing our resources into making happy mice. Like humans, mice take pleasure in sex, play, physical activity, exploring their environment and so on. Again, it may be that the pleasure humans take in these activities is greater than that of mice. As we saw above however, we can feed at least 600 mice for the same investment it takes to feed a single human. For human pleasure to outweigh mouse pleasure then we would have to believe that the pleasure experienced by a human being engaged in one of the above activities is at least 600 times greater than the pleasure of a mouse engaged in the same kind of activity.

The above examples help clarify the notion of EWP and suggest that mice may be better converters than humans. It would be useful to have harder data on the welfare realizing capacities of mice and of the costs of keeping them alive and happy. New research by the organization Rethink Priorities (2022) may help shed some light on the former. The 'Moral Weight Project' is an effort to get a better understanding of the capacities for welfare of several commonly farmed animals. The study assumes hedonism, then based on several hypotheses about the function of valenced experience they identified 96 measurable traits that can be used as proxies for welfare capacity. By surveying the literature on the behavioral and cognitive capacities of the selected animals they were able to estimate figures for the welfare ranges of each, where welfare range is the difference between the most intense positively valenced state the individual can realize and the most intense negatively valenced one. These ranges were then compared to the human range. The results cannot be expected to be precise, but they are suggestive. The welfare ranges of mice were not estimated. However, the welfare range of pigs was estimated to be slightly over half that of humans, and the range for chickens was one third that of humans.

As a reasonable figure, based on the findings of *Rethink Priorities* for pigs and chickens, I will assume that mice have a welfare range that is one quarter that of humans. That is, mice are capable of a range of positively and negatively valenced states up to 25% of the intensity of human states at any given moment. First, mice are social and intelligent animals, and they display many of the traits that the study uses as proxies for welfare capacity, such as play behavior. Second, one tentative conclusion of the study is that, for all the vertebrates studied, none has a welfare range that is more than twice that of any other. Since the study also included carp and salmon, we would have to believe that mice have a lower welfare range than these animals in order to have a range lower than 0.25. Of course, this may be the case, but it is very counterintuitive. Finally, mice are more closely related to humans than chickens *and* pigs. Our last common ancestor with mice was around 70 million years ago, compared with 90 million for pigs, and over 300 million for chickens. None of these

[13]One can purchase enough food to feed a mouse for over a year for about £7.

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

considerations are decisive, but they do suggest a relatively high welfare range for mice. Until further research is done, I will continue to assume a welfare range for mice of 0.25 that of humans.[14]

Based on the trivial cost of feeding and housing mice, and the reasonable assumption that they can generate 25% as much positive welfare as a human being at a given time, it seems that a population engineer who was impartially concerned with creating as much positive welfare as possible might try to move us toward a population composed more of mice than of humans and other species. This is because they have a higher EWP than humans do, and in a world with finite resources we have reason to take the most efficient route to our goals. We can maximize welfare more efficiently by making happy mice rather than making happy people. If the arguments in this section are right, then we have reason to aim for a very large population of happy, domesticated mice, with the minimum human population necessary to maintain it. Ultimately, we may not need humans at all, if we can be replaced with artificially intelligent machines who can be given the task of looking after the mice for us. In fact, if we develop digital sentience, it may be possible to create positive welfare even more efficiently by creating digital minds rather than biological ones, if digital sentience is more resource efficient than biological sentience. This might seem like a way out of the problem I have raised here – if we could create human-like digital minds that are more efficient than animal minds, then we might be able to maximize total welfare without having to replace ourselves with simpler animals. Unfortunately, this does not solve the problem – since animal brains are so much simpler than human ones, they will likely be much easier to replicate digitally, and consequently digital animal minds would have higher EWPs than digital human ones. Our population engineer would have reason not to create digital humans, but digital animals instead.[15]

3.2 What's bad about mouse world?

The idea that animal-Z might be more valuable than an A world of extremely happy humans seems repugnant; the idea that we might have moral reason to try to move closer to such a world seems even worse. Yet if we accept the assumptions I have made above, it looks like we do have such reasons. It is worth spelling out exactly what seems so repugnant about that conclusion, and what would be missing from animal-Z.

The first thing that seems bad about animal-Z is the very low welfare in each life. The welfare range of a mouse may be about 25% that of a human being, meaning the maximum positive welfare they can experience at a given time is about one quarter as intense as that available to a human. Even worse, since each mouse only lives for a couple of years, while humans might live to 80, the total lifetime welfare of each mouse must be divided again by 40. Taking these figures literally then, a mouse in animal-Z would have a lifetime welfare level about 160 times lower than a human in any possible human-Z world. Although in thought experiments we can imagine humans having such low lifetime wellbeing levels, in practice it seems hardly possible to have a viable population of human beings at such a low level of lifetime wellbeing. Average welfare in animal-Z is much lower than that in any possible human world approaching Z.

[14]Bob Fischer at Rethink Priorities is currently researching the welfare ranges of mice. In personal correspondence he said he would put the welfare range of mice at around the same level as that of chickens, that is about 0.33 that of a human. However, he stressed that he has low confidence in this estimate.

[15]See O'Brien (2024, p. 11) and Fischer (unpublished work).

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

The second bad thing about animal-Z is the lack of all the higher pleasures and best things in life. As we saw, these things might also be absent from a human-Z world, at least in theory. In practice though it is hard to see why any realistic human-Z-like world would lack all these good things, especially since many of the higher pleasures are cheap to provide (think of digital copies of all of Shakespeare's work, recordings of Mozart's music etc.). Animal-Z by contrast would certainly lack all of these good things. It should be stressed here that these are not purely theoretical points, rather they are made more pressing by the fact that we could bring about such a state of affairs. As I have tried to emphasize, the difference between a human-Z and an animal one may be relatively trivial when we consider those worlds as thought experiments. Without the constraints of reality, we can imagine humans with lifetime welfare levels as low as those of mice, humans who, for some reason, are unable to enjoy any of the higher pleasures characteristic of human life and so on. A realistic human-Z world is unlikely to be quite so bad – it is likely that there is a practical 'welfare floor' below which it would not be possible to sustain a large population of humans, and there is no reason to think that all of the best things in life would be lacking from any realistic human-Z world – whereas a realistic animal-Z would be very bad indeed in terms of average welfare and loss of the best things in life.

The third thing lacking from animal-Z is species diversity. The reason for this is simple. Our population engineer's strategy for maximizing welfare given her limited resources was to discover the species with the highest EWP and to try to increase their population as much as possible. If successful, this would involve taking resources away from less efficient species and allocating them to the more efficient ones. Related to this is a lack of wild animals generally. Wild animals are a waste of resources in two ways. First, it is plausible that most of them have bad lives.[16] Second, even those that have positive lives, like elephants and chimps, are much less efficient at converting resources into wellbeing than small animals. While a committed total utilitarian might be happy to sacrifice biodiversity and wildness to attain more positive welfare, most people would not be so blasé about this loss. It is important to note here that I am *not* claiming that we should never sacrifice biodiversity in order to reduce wild animal suffering. We may have strong moral reason to try to reduce animal suffering even if this comes at the cost of reducing biodiversity. I am making the much weaker claim that it is very counterintuitive that we should totally sacrifice biodiversity in order to maximize positive welfare. If Totalism implies that we should completely sacrifice the impersonal value of biodiversity in order to further increase positive welfare, then this may make it more unappealing as an account of the value of populations.

In practice, a perfectly implemented animal-Z world seems unlikely, even if it is possible. However, the Normative Problem it raises does not go away entirely just because animal-Z is unlikely to be fully implemented. Creating happy mice is so cost effective that it ought to be taken seriously as a way of making the world better, at least by utilitarians, and it is surely implausible that they have all things considered reasons to create colonies of happy mice.[17] Secondly it suggests that we should not regret losses in

[16]See Ng (1995), Horta (2010), Johannsen (2021), O'Brien (2022) and Faria (2023). For criticism of this view see Browning and Veit (2023) Gardner (2022) and Groff and Ng (2019).

[17]For an argument that egalitarianism might require a massive shift of resources from humans to mice, since mice are so badly off compared to humans, see Vallentyne (2005). Matheny and Chan (2005) argue against the 'logic of the larder', that is, the idea that animal agriculture benefits animals as it causes them to exist with lives worth living. Part of their argument is that, if we take the logic of the larder seriously, and

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

wilderness and losses in species, insofar as these bring us closer to the ideal state of affairs. The idea that we have moral reason to try to create a world which has (1) such low average welfare (2) lacks the best things in life (3) is severely lacking in biodiversity and (4) lacks wild animals living free from human control is enough to make us doubt that a Totalist welfarist axiology is sufficient for thinking about the value of populations that differ in species composition.

Section 4

4.1 Population engineering 2: human world

In this section I will assume that there is something special about the constituents of welfare available to typical humans which means that human beings are the species with the highest EWP. This could be because the higher pleasures available to human beings are lexically superior to those lower pleasures available to simpler animals as Mill claimed. This kind of strong lexicality implies that a single human life containing the higher pleasures would be better than any number of lives without them. Some have tried to use more sophisticated forms of lexicality to avoid the RC without entailing this equally implausible conclusion. Nebel (2022) has developed a view of lexical thresholds which allow him to avoid this conclusion, and Parfit (2016) has suggested a form of perfectionism according to which the 'best things in life' are imprecisely lexically superior to lower quality components of welfare such that any loss of these best things cannot be compensated for by increases in the amount of lower quality constituents of welfare.

Since I am concerned with the Normative Problem rather than the axiological problem, I need not assume any form of lexicality. If my concern was with axiology, then I would have to provide an account of lexicality that gives the right answers in hypothetical cases involving the existence of impossibly vast numbers of nonhuman animals and much smaller human populations. Claiming the superiority of the smaller human population in imaginary cases requires something like lexicality, as we can simply keep increasing the numbers of animals in the imagined animal population. Since I am concerned with the Normative Problem I can simply assume that those higher human pleasures produce vastly, but finitely, more welfare than lower ones, as there is of course a limit to how resource efficient simpler animals can be compared to humans. That is, there is a class of distinctively human activities that contributes so much to welfare compared to those simpler pleasures also available to nonhuman animals that it would be practically impossible to create enough simple animals living good lives that would produce more welfare value than a smaller population of humans engaged in those activities. I assume that this class of distinctively human activities contains things like those identified by Mill as the higher pleasures, and by Parfit as the best things in life, such as doing philosophy, listening to excellent music and so on.

Again, our population engineer wishes to maximize the welfare value of the global population given finite resources, and once again there is one species which has a higher EWP than any other, humans. How might she proceed? A rational program would be to increase the human population as much as possible, insofar as this is compatible with humans continuing to engage in these highly welfare efficient activities. This would

want to use our money to create happy animals, there are more efficient ways of doing so than farming animals. Since mice have very low resource requirements, we ought to renounce animal agriculture and "adopt a vegan (vegetarian) diet and invest savings in colonies of mice" (p. 589).

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

presumably involve maximizing the living space available to human beings – insofar as it is possible, she would want to 'pave over' wilderness to make room for more humans. This would mean reducing wild animal populations significantly. It might also involve significant reductions to the population of domesticated animals. Since most domesticated animals are used for food, and since plant-based diets are far more resource efficient, she would have reason to eliminate wasteful practices of animal agriculture to maximize the human population. The end goal of her efforts would be a vast population composed almost exclusively of human beings, each of whom spends their lives enjoying the most resource efficient of the human welfare producing activities.

4.2 What's bad about human world?

Appealing to a more complicated conception of welfare according to which humans are vastly better at generating welfare than simpler animals has allowed us to avoid the conclusion that animal-Z is more valuable than a human A population in practice. In some ways the resulting human world is quite appealing – there is a large human population enjoying a very high total and average level of wellbeing, and not only are the 'best things in life' not missing, they are central to the lives of each person. There are some potentially important things missing from this world too, however. Like mouse world, this world is severely lacking in diversity of animal life and of wilderness generally. The reason for this lack is the same in both worlds – in each world there is one species with a higher EWP than all others, and, given the fact of finite resources, in order to maximize the welfare of the population our population engineer ought to try to shift the species composition of the population so that it contains as many of these efficient animals as possible. Other, less efficient animals ultimately become a waste of resources. Wilderness too is a waste – if most wild animals have bad lives it is better to replace these wild spaces with human controlled ones. Even if animals in the wild live marginally net positive lives, it is unlikely that uncontrolled wilderness would be the optimal way to produce the highest total welfare.

In practice, our population engineer is unlikely to be able to fully implement this human world either. Still, some unpalatable normative conclusions remain even if full realization proves impossible. First, if the arguments in this section are correct then we have less reason to worry about human caused extinctions, and the constant encroachment of human beings into formerly wild areas. In fact, we have reason to continue such activities so long as continued human expansion does not threaten our long-term ability to maximize our population. We should see animals, even those with net positive lives, not as wondrous fellow creatures sharing the world with us, but as regrettable axiological cul de sacs, inefficient resource sinks that prevent us from maximizing total positive welfare. We might wish to care for those animals who exist, while arranging for their gradual extinction. If we wish to maximize welfare, then the long-term future ought to be a human one.

Conclusion

In this paper I have investigated some of the consequences of including animals in our thinking about the value of populations. I have argued that the RC can easily be extended to include animals, resulting in animal-Z. While the human-focused RC may be a merely theoretical problem, including animals raises practical implications that few would be willing to accept, and this puts pressure on Totalists. For those

https://doi.org/10.1017/S0953820824000098 Published online by Cambridge University Press

who are undecided, Totalism's unappealing normative implications once animals are included in the analysis may count against it as an adequate population axiology.

Committed Totalists might accept the normative implications I have outlined, despite their unpalatability. Which kind of population best maximizes total welfare then turns out to depend essentially on what conception of welfare we adopt. If we adopt simple hedonism then the fact that some animals can achieve levels of welfare comparable to those attained by humans, but in a much more resource efficient way, suggests that the ideal population would be one composed of relatively simple animals such as mice. If on the other hand we think that human welfare is far superior to that of other animals, then the ideal population may be one composed entirely of human beings engaged in typically human activities of the sort celebrated by Mill and Parfit. I suggested that neither of these populations is intuitively very appealing, and that both are severely lacking in species diversity and wild animals. If Totalism really implies that one of these two populations would be ideal, then this is a counterintuitive implication of that view. For those of us who, like me, find the two worlds I have described above unappealing, it may be the case that we need to adopt a richer population axiology than Totalism if we are to accommodate animals in the populations analyzed.

Acknowledgements. I am very grateful to Guy Kahane, Jeff McMahan, Beatrice Marchegiani, and two anonymous reviewers for this journal, each of whom gave detailed feedback on multiple drafts of this paper. For further comments and discussion, I would like to thank Mattia Cecchinato, Lorenzo Elijah, Kida Lin, Katie Prosser, Diego Expósito Teixeira, Patrick Williamson, and attendees at workshops in Sheffield University, Lingnan University, and the MIMIR Center for Long Term Futures Research.

Competing interests. None.

References

Beckstead, Nick. 2013. *On the Overwhelming Importance of Shaping the Far Future*. PhD thesis, Rutgers University, NJ.

Bentham, Jeremy. 1789. *An Introduction to the Principles of Morals and Legislation*. Oxford: Clarendon Press.

Browning, Heather. 2023. Welfare Comparisons Within and Across Species. *Philosophical Studies*, **180**: 529–51.

Browning, Heather and Veit, Walter. 2023. Positive Wild Animal Welfare. *Biology and Philosophy*, **38**(14): 1–19.

de Lazari-Radek, Katarzyna and Singer, Peter. 2014. *The Point of View of the Universe: Sidgwick and Contemporary Ethics*. New York: Oxford University Press.

Faria, Catia. 2023. *Animal Ethics in the Wild: Wild Animal Suffering and Intervention in Nature*. Cambridge: Cambridge University Press.

Fish Count. 2019a. Numbers of Farmed Decapod Crustaceans. Available at: http://fishcount.org.uk/fish-count-estimates-2/numbers-of-farmed-decapod-crustaceans [accessed 21/07/23].

Fish Count. 2019b. Numbers of Farmed Fish Slaughtered Each Year. Available at: http://fishcount.org.uk/fish-count-estimates-2/numbers-of-farmed-fish-slaughtered-each-year [accessed 21/07/23].

Gardner, Molly. 2022. Suffering and Meaning in the Lives of Wild Animals. *Midwest Studies in Philosophy*, **46**: 355–71.

Greaves, Hilary. 2017. Population Axiology. *Philosophy Compass*, **12**(11): e12442.

Groff, Zach, and Ng, Yew-Kwang. 2019. Does Suffering Dominate Enjoyment in the Animal Kingdom? An Update to Welfare Biology. *Biology and Philosophy*, **34**(40): 1–16.

Healey, Richard and Pepper, Angie. 2021. Interspecies Justice: Agency, Self-Determination, and Assent. *Philosophical Studies*, **178**: 1223–43.

Horta, Oscar. 2010. Debunking the Idyllic View of Natural Processes: Population Dynamics and Suffering in the Wild. *Telos*, **17**: 73–88.

Huemer, Michael. 2008. In Defence of Repugnance. *Mind*, **117**(468): 899–933.
Johannsen, Kyle. 2017. Animal Rights and the Problem of r-Strategists. *Ethical Theory and Moral Practice*, **20**: 333–45.
Johannsen, Kyle. 2021. *Wild Animal Ethics: The Moral and Political Problem of Wild Animal Suffering*. New York: Routledge.
Matheny, Gaverick and Chan, Kai M. A. 2005. Human Diets and Animal Welfare: The Illogic of the Larder. *Journal of Agricultural and Environmental Ethics*, **18**: 579–94.
Mctaggart, John M. E. 1927. *The Nature of Existence (Volume II)*. Cambridge: Cambridge University Press.
Nebel, Jacob. 2022. Totalism Without Repugnance. In J. McMahan, T. Campbell, J. Goodrich, & K. Ramakrishnan (Eds.), *Ethics and Existence: The Legacy of Derek Parfit*. Oxford: Oxford University Press, 200–31.
Ng, Yew-Kwang. 1995. Towards Welfare Biology: Evolutionary Economics of Animal Consciousness and Suffering. *Biology and Philosophy*, **10**: 255–85.
O'Brien, Gary David. 2022. Directed Panspermia, Wild Animal Suffering, and the Ethics of World-Creation. *Journal of Applied Philosophy*, **39**(1): 87–102.
O'Brien, Gary David. 2024. The Case for Animal-Inclusive Longtermism. *Journal of Moral Philosophy*, (published online ahead of print 2024). https://doi.org/10.1163/17455243-20234296.
Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
Parfit, Derek. 1986. Overpopulation and the Quality of Life. In P. Singer (Ed.), *Applied Ethics*. Oxford: Oxford University Press, 145–64.
Parfit, Derek. 2016. Can We Avoid the Repugnant Conclusion? *Theoria*, **82**(2): 110–27.
Plant, Michael. 2022. The Meat Eater Problem. *Journal of Controversial Ideas*, **2**(2): 1–21.
Rethink Priorities. 2022. The Welfare Range Table. Available at: https://rethinkpriorities.org/publications/the-welfare-range-table [accessed 14/09/23].
Sebo, Jeff. 2023. The Rebugnant Conclusion: Utilitarianism, Insects, Microbes, and AI Systems. *Ethics, Policy and Environment*, **26**(2): 249–64.
Sidgwick, Henry. 1907. *The Methods of Ethics*. London: Macmillan.
Tännsjö, Torbjörn. 2002. Why We Ought to Accept the Repugnant Conclusion. *Utilitas*, **14**(3): 339–59.
Thornton, Alex. 2019. This Is How Many Animals We Eat Each Year. *World Economic Forum*. Available at: https://www.weforum.org/agenda/2019/02/chart-of-the-day-this-is-how-many-animals-we-eat-each-year/ [accessed 21/07/23].
UNEP. 2012. *One Planet, How Many People? A Review of Earth's Carrying Capacity*. Available at: https://na.unep.net/geas/archive/pdfs/geas_jun_12_carrying_capacity.pdf [accessed 21/07/23].
Vallentyne, Peter. 2005. Of Mice and Men: Equality and Animals. *The Journal of Ethics*, **9**(3/4): 403–33.
Williamson, Patrick. 2021. A New Argument Against Critical-Level Utilitarianism. *Utilitas*, **33**: 399–416.
Zuber, Stéphane *et al.* 2021. What Should We Agree On About the Repugnant Conclusion? *Utilitas*, **33**: 379–83.

Cite this article: O'Brien, Gary David. 2024. Totalism, Animals, and the Repugnant Conclusion. *Utilitas* **36**, 211–229. https://doi.org/10.1017/S0953820824000098