

Longtermism and Animals

Heather Browning and Walter Veit

1 Introduction

When deciding how to act in a world of limited resources, we must have methods or guidelines for the prioritisation of actions that lead to the best outcomes, or at least avoid the worst. Roughly speaking, longtermism is the ethical doctrine that the rightness of our actions is primarily determined by their effects in the long-term future. When making ethical decisions, it is not only the present or near future that matters, but all future individuals and events. As there is, in expectation, much more value in the long-term future than the present or short-term future, the best actions will thus be those that have the best effects in the long-term future, shifting our focus of attention towards interventions that provide such long-term future benefits (Beckstead 2019; Greaves and MacAskill 2019; MacAskill 2022).

It's highly likely there will be far more people in the future than there are in the present, or the past. Not only is human history relatively young, in evolutionary terms, but there is also good reason to think that future technology would allow for a larger number of humans at any one time. The number of potential future humans has been estimated at the low end at 1 quadrillion (10^{15}), which is 100,000 times more than currently exist—and this is only with taking numbers at-a-time as remaining fairly stable (Greaves and MacAskill 2019). The number only increases when considering future technologies that allow for larger population sizes, particularly those that allow for human migration beyond Earth, thus opening up the possibility of population expansion of many orders of magnitude (Bostrom 2003). Additionally, ongoing scientific and technological developments mean that these people are likely to have a higher quality of life than our own (Beckstead 2019). Because of the overwhelming numbers of future people, the argument goes, we are morally required to focus on ensuring the long-term future goes well for these people. This means, when choosing between actions, we base our calculations on the effects in the long-term future (~1000+ years). Calculating the expected long-term value of present actions is thus the primary activity of those aiming to operate under a longtermist framework.

There is a growing literature on the strengths and weaknesses of a longtermist viewpoint, particularly regarding its tractability and underlying axiological commitments (Greaves and MacAskill 2019; Tarsney 2020; Thorstad and Mogensen 2020; John and MacAskill 2021; Mogensen n.d.), and it is not our aim here to assess its merits. Here instead we wish to focus on what has been too often overlooked in many discussions of longtermism—the consideration of non-human animals. Almost all the current writing on the topic references humans, and the proposed and debated interventions are also those which benefit human populations, such as reduction of existential risk (Bostrom 2003; 2013; Ord 2020; MacAskill 2022) and promotion of technologies that enhance our capacities to expand our population in future, particularly on other planets (Bostrom 2003).

While these are important concerns, also important are those that consider the long-term future of non-human animals. Although almost all moral theories accept that non-human animals are important sources of moral value, the focus in longtermism has so far been almost exclusively human. For instance, *What We Owe the Future* (MacAskill 2022), a book that has arguably introduced longtermist thinking into the public sphere, devotes only a single subsection of one chapter to considerations of non-human animals. Tarsney (2020) recognises this omission within longtermist research, but comments that: '(1) The sign and magnitude of the effects of paradigmatic longtermist interventions on the welfare of non-human animals (or their far-future counter-parts) are very unclear. (2) Dropping this simplification seems unlikely to change our quantitative results by more than 1–2 orders of magnitude (though this is far from obvious), and so unlikely to affect our qualitative conclusions' (36). Here, we strongly disagree with both of these contentions—that we cannot know what the long-term future would be like for animals (in a way that differs from our uncertainty for humans), and that the orders of magnitude of our results would be largely unaffected. In this chapter, we will argue that the interests of animals are just as important to consider as those of humans, and in our deliberations over best actions, animals should be given much more consideration than they currently are, but that this is a research area in longtermism that is currently neglected.

2 Why animals should count

2.1 Numbers

There are vastly more animals on the planet than there are humans. Even if we only count vertebrates (as these may be the only animals we can currently reliably identify as sentient and thus capable of morally relevant states of pleasure and suffering), there are over 100,000 animals for every human (estimated 10^{11} land vertebrates and 10^{15} ocean to 10^{10} humans) (see Bar-On, Phillips, and Milo 2018). While many wild populations are shrinking, numbers of domesticated animals, particularly in agriculture, are rising. Every year, somewhere around 90 billion fishes, 70 billion chickens, 300 million cows, 1 billion sheep and goats, and 1.5 billion pigs are raised and killed for food,¹ and an additional 1–3 trillion fish taken from the oceans.² This is more *annually* than the number of humans that have ever existed. These numbers are hard to even conceptualise, and yet, they would grow even more if we were to consider the human impacts on invertebrates. If current production and consumption habits were to remain unchanged, it is clear that there would continue to be exponentially more animals than there are humans. Thus, if the long-term future matters because of the large number of humans it contains, it should equally matter for the even larger number of animals. Concern for animal interests will be a high priority simply because there are *just so many of them*.

One way to resist this could be to argue that although there are many animals, they should count for less in our calculations of expected value (e.g. MacAskill 2022). We will take it here as uncontroversial that animal welfare should count for something under most

¹ Numbers from Šimčíkas (2020)

² Numbers from fishcount.org.uk (2019)

conceptions of value. This does not require equal consideration of the interests of humans and non-human animals—we can accept that species membership may change the strength of interests, or the total level of pleasure or suffering experienced, such that animals will be weighted differently in calculations to humans. This demonstrates the urgent need for interspecies comparisons of welfare as it is only through performing such comparisons that we can make the necessary calculations to determine in which cases animal or human considerations will dominate. The interspecies comparison problem is a complex one (see Browning (2023) for some discussion) and research into it should form a priority for any longtermist research programme. However, unless we assign only an extremely (and arguably, implausibly) low weighting to animals, their sheer numbers mean that they are still likely to dominate humans by several orders of magnitude.³ The same failure of comprehension that longtermists try to combat regarding the number and importance of future humans is seemingly still at play when considering the number of current and future animals.

One could also counter that although there are undoubtedly currently more animals than humans, that this won't be the case in the future. For example, we might think that the societal shifts we can currently see in the rise of veganism mean that factory farming will be phased out at some point in the medium-term future, so these animals will not exist in the long term. We will address this concern further when we talk about this intervention, but here we will just note that it is not at all obvious that this will actually be the case, without more action than is currently being taken. Or we might think that the number of wild animals will decrease, as we head into another potential mass extinction event. However, even if such an event does occur, it will be a reduction in species diversity, not necessarily a reduction in total numbers—those animal species that do well in human-altered environments (such as urban pests) are likely to continue to thrive. For example, climate change could alter the distributions of species such that insect populations are able to expand further north and south, increasing the numbers of these animals even if some larger animals decline (Sebo 2022).

Lastly, we might think that when humans move out to colonise other planets we will do so without other animal species, and thus our future growth will vastly outstrip theirs. In particular, if we think that it is the small probability of this large explosion in human population size that creates most of the expected value of the far future (e.g. Tarsney 2020), then this will be the most important determination of whether or not animals will also count. There is no simple reply to this. The details will depend a lot on the specific methods used in interstellar expansion, which would currently seem to be an open question, dependent on future technology. However, there are a couple of ways in which animals would remain an important source of value in terms of their numbers. The first is if we continue to use agricultural animals as a means of sourcing easy protein, as may be the case when setting up new settlements. The second is if we colonise by way of terraforming, creating planetary ecosystems to support human and other forms of life. Even if the number of animals taken to begin such processes is small, creation of any flourishing ecosystem is going to very quickly lead to a large number of animals.

It is also possible that the future will not be dominated by either humans or non-human animals but digital beings—sentient AIs. In the end, there is a lot of uncertainty here and

³ Current attempts to weight based on neuron count (e.g. MacAskill 2022) are unconvincing (see Shriver 2022).

unless we are quite sure of these alternative outcomes, we still have reason to believe that there will be very high numbers of animals in the future.

2.2 Suffering

As well as there being lots of animals (both now, and expected in the future), many of these animals will have bad lives. In the words of Beckstead there are: ‘an astronomical number of expected future beings with lives that are suboptimal, and a future whose trajectory is potentially influenceable’ (Beckstead 2019: 92). Though he was talking about pessimistic estimates of the lives of future humans, the same applies even more strongly for animals. There is thus a great amount of future suffering that we can potentially prevent.

From the numbers we presented above, we can see that almost 75% of land vertebrates live in agricultural systems. These systems are well known for the suffering caused to the animals (Harrison 1964; Singer 1975; Gruen 2011). Most broiler chickens spend their lives in windowless sheds with under one square foot per bird; their beaks are trimmed using hot blades to decrease the aggression brought on by the crowded conditions. They frequently suffer leg deformities and lameness from ongoing selective breeding for rapid growth. Sows used for breeding are often kept in tiny stalls in which they are unable even to turn around, with few cognitive or behavioural challenges/opportunities and no access to nesting materials to fulfil their strong drive for nest-building. For many, if not most, of these animals, there are almost certainly ongoing negative experiences and few opportunities for positive experiences such that their lives are highly likely to contain more suffering than pleasure. If current agricultural practices were to continue like this into the future, there would be ongoing suffering at a large scale. Again, one may counter that we should not expect high levels of future animal suffering simply based on current circumstances. If factory farming is going to end, or if conditions are going to vastly improve, then we will not have future suffering of food animals. As we will argue in what follows, even if this is true we may still see huge benefit in speeding up the trajectory.

Many wild animals also suffer. Many writers argue that, in fact, suffering dominates in nature (Ng 1995; Horta 2010; Tomasik 2015; Iglesias 2018). This is in part attributed to the general causes of suffering, such as injury, disease, starvation, and predation. However, it is also considered to be an effect of the life history of many wild animals—the ‘r-selected’ species that produce a large number of small or ‘cheap’ offspring, of which only a few live to maturity. The large numbers that instead perish are considered to have lives almost completely composed of suffering (from whatever processes kill them), with few if any opportunities for pleasure. Given the large numbers of such individuals, it is then taken to be the case that there is an overwhelming prevalence of suffering over pleasure. Though we think there are reasons to doubt that animal suffering in the wild outweighs positive experience (Browning and Veit 2023), it is obvious that it is still widespread. Overall, not only are there lots of animals, but they potentially have lives containing a lot of suffering, and that we can change for the better. Animal suffering is a major, if not *the* major, source of current disvalue, and plausibly so too in the long-term future. It should thus be accounted accordingly.

3 Potential interventions

We have argued here that it is important to include animals in calculations about which actions we should prioritise for the long term; which, due to the numbers and degree of suffering involved, is likely to lead to giving priority to animal-based actions in many cases. We do not rule out that in some (or even most) cases, the calculations will still favour human-centred interventions, for a range of reasons such as tractability or differential moral weight, but this should not be taken for granted without further investigation. Instead, relevant animal-based actions should also be assessed and compared. How, then, can this be done? In this section we will survey a number of different potential interventions that may improve the long-term future for animals. We don't make any strong claims about which would in fact be the best options to pursue, but discuss what we take to be some of the more promising avenues for further investigation.

The important categories for action in shaping the far future can be divided into 'proximate benefits', speeding up development, and trajectory changes (Beckstead 2019). Proximate benefits are the more predictable, short-term benefits of action. Speeding up development refers to pushing developments that could improve future quality of life to earlier in the timeline, such that their benefits will be felt for longer. Trajectory changes are arguably the most impactful, and involve shifting the direction of the world's development, such that we end up with a different kind of future than we otherwise would have had—an example of this being the abolition of slavery (MacAskill 2022). A current example relevant to animals may be the rise of global aquaculture—this is an industry that is still young, and the structures and regulations we set in place now may have long-reaching effects in terms of how the industry develops (Franks, Ewell, and Jacquet 2021).

One potential objection is that actions to improve human welfare may be the right priority right *now*, as ensuring the welfare of humans is also the best way to create a future in which animals are taken care of (Sebo 2022). Improving our social, economic, and political systems can help empower future generations and create space for developing a capacity for and desire to help animals. Not until our own needs are met can we perhaps then turn to assisting others. This also relates to a possibility that if most of the value in the far future will actually be realised by sentient AIs, created by humans, then we should be prioritising actions to protect humans and ensure the development and spread of such AIs. While we see value in this objection, it is not one that can just be asserted *a priori*. It may very well turn out to be the case as a result of our calculations, when we place both human and non-human animal wellbeing into the calculus. Importantly though, this decision needs to be made after making an assessment including all these factors, and with comparison to the set of possible alternative actions focusing more directly on animals. We are not claiming that these assessments would end up showing we should prioritise interventions to help animals, but merely that without including consideration of animal interests, we couldn't know for sure. We should also be wary of potential motivated reasoning toward conclusions that support our own self-interested preferences, and making sure to include reasoned considerations of non-human animal interests will help prevent this.

There are of course many different possible actions that could help improve the situation for animals in the long-term future, but here we will outline a few that are likely to be beneficial and are worthy of further investigation. They can be grouped into two categories—those

that change the number of future animals, and those that change the quality of life of future animals (changing the size of the future and changing its sign). That is, we should try to ensure there are lots of future animals if we predict their lives to be good, and few if we predict them to be bad. Additionally, we should work to try to improve expected quality of life such that all the animals who will exist will have lives of the highest positive welfare we can achieve. We take it that in response to observed suffering it is preferable to act to reduce the sources of suffering rather than the number of bearers of suffering (Višak 2017), wherever the former is possible. But which of these interventions different actors prefer will depend strongly on their ethical and axiological commitments.

3.1 Changing the number of animals

There are two ways in which we can beneficially change the size of the future regarding animals—one is in reducing the number of animals if we expect them to have bad lives, and the other is increasing the number of animals we expect to have good lives. This may be a complex question to answer in practice, as the differences in animal cognitive sophistication, lifestyles, and evolutionary history will influence their overall lifetime welfare balance—for example, prey animals may experience more fear from predator presence, while predators may be more stressed by the demands for successfully finding and hunting prey. However, in thinking about setting up a long-term future that contains few suffering animals and abundant happy animals it is important to think about which animals will have good or bad lives. This differs from changing the quality of animal lives from negative to positive as what we're considering is decisions about whether or not to bring animals into existence rather than how to make their lives better.

For the first—reducing the number of unhappy animals—one potentially important intervention is ending factory farming. Though the numbers involved are lower than for wild animals, the suffering is arguably higher—with most animals probably having strongly net-negative lives—and this is a more obviously tractable intervention than many of those discussed for wild animals. Thus, ceasing to bring animals into these situations would be a significant change to overall value, and a long-term future in which such practices no longer exist will be a far better one than if they do. For instance, widespread adoption of a vegan diet would lead to fewer animals used, and thus fewer numbers over time. Development of *in vitro* 'clean meat' products is one possible path to this end (Anomaly et al. 2024), as are general advocacy movements to increase veganism. In general, intensive farming is benefitted through direct subsidies and by externalising the costs of harms to health, environment, and animal welfare; simply altering these would make the industry far less economically viable (Sebo 2022). This is an example of where longtermist and short-termist goals align—reducing the number of suffering animals *now* and preventing far more being created in the future. Additionally, if, as John and Sebo (2020) argue, the existence of animal agriculture maintains human attitudes toward animals that hinder moral circle expansion, then elimination of this practice is crucial to ensuring the ongoing wellbeing of animals in the long-term future.

This could be considered a version of 'speeding up progress', if we think that animal agriculture will eventually die off, but that the sooner we reduce it, the more animals will be saved from coming into existence in a life of suffering. Given the large numbers and

suffering involved every year that intensive farming persists, any action we can take to bring this sooner will still represent a large gain. We can also here include possible additional benefits, such as reducing risks of future pandemics, and slowing down climate change. As will be discussed further on, it might also be considered as a movement toward a better attractor state. If we think that at some point the dietary preferences of humans will become fairly fixed in one state or the other, then pushing toward the higher-value state would ensure a better future. As we have already mentioned, we may resist this as a longtermist priority if we think that we are already on this path, such that factory farming is likely to end in the short- to medium-term future. Evidence for this could be seen in the increased adoption of a vegan diet (Russell 2023), and growing concern for the welfare of farmed animals; but on the other side we can see numbers of intensively farmed animals still continually increasing (Torrella 2021). While the population share of vegans increases, so too does the total number of humans who aren't. A lot here depends on where we see the current trend heading, and whether or not intervention now is needed to ensure this state in the long-term future. If we are at all uncertain about this trajectory, actions to ensure we bring about the more positive outcome would have a high expected value.

Another way in which we could reduce the number of suffering animals could be in reducing the number of wild animals, or at least those of the types we take to have lives predominantly composed of suffering. This seems to be the view taken by some writers (Tomasik 2017), who follow the 'logic of the logger' (John and Sebo 2020) in arguing that reduction of suffering entails habitat destruction, to decrease the number of animals. Ensuring a future with fewer or no suffering animals will increase its expected value. It is an open question as to how much of a current priority this should be, based on which specific actions now are likely to have uniquely strong effects on the numbers of wild animals in the far future. We take this to only be desirable if we are unable to instead intervene and improve the lives of these animals, an option we will discuss in the next section. While the former would reduce the amount of disvalue, the latter would also increase the amount of value, which will bring greater overall benefit.

The other way in which we can positively impact the size of the future is in ensuring there are large numbers of animals with positive welfare. The far future will have far greater total value with a large number of happy animals in existence than if suffering animals are simply absent. If it is the case that the existence of more net-positive lives is worthwhile, then we should be looking at ways to maximise the number of happy animals. One version of this would be mitigating extinction risk, at least for species with good lives. Like the mitigation of human extinction risk, this would allow for a future filled with much larger numbers of happy beings. Many of the efforts to mitigate extinction risk will align with those used for humans (e.g. addressing climate change, reducing the chance of meteor collision), but there will be some unique to animals. For many of these actions there will be complex trade-off calculations necessary, as resource distribution considerations require that increasing numbers of some species will place limits on others.

Prioritising the creation of animals who would have good lives would involve determining which animals may be capable of the most pleasure, and the conditions under which they should be kept to realise it, then investing resources in their creation and management. Understanding the relative sentience of different creatures, as well as allowing us to assess their relative suffering, will also give us guidance on what sorts of creatures we should be creating—which provide the most potential 'welfare per unit', so to speak.

If, for example, dogs are capable of as much pleasure as humans, but it is much simpler to provide them with what they need to achieve it, this gives us reason to promote the future numbers of dogs over those of humans. Depending on the empirical facts about relative wellbeing, it may even turn out that it is worth sacrificing potential numbers of humans in order to make this happen. This would also require understanding of the longer-term effects and side-effects—both positive and negative—of population expansions to judge the relative benefit.

3.2 Improving the lives of animals

As well as changing the size of the future we can also aim to change its sign—that is, to reduce suffering and increase pleasure for those animals who will exist. This could be done for agricultural animals and/or wild animals, or for the additional human-created animals described in the previous section. Reduction or elimination of animal agriculture is important where we think that animals in these conditions have net-negative lives, which is highly plausible for most modern practices. There is, however, also the possibility of changing farming practices such that animals have net-positive lives. In this case, the so-called ‘logic of the larder’ (John and Sebo 2020) would then advocate their creation and consumption, because the creation of positive lives is an overall good and should be encouraged. However, this is unlikely to have the expected benefits. If what we wanted to do was create the greatest number of happy animal lives, it is doubtful that agriculture would be the most cost-effective way to do so. We could, for example, raise large colonies of happy mice for far less money than the agricultural industry takes to sustain, as well as freeing up cropland currently used to feed agricultural animals, providing habitat for more wild animals (Matheny and Chan 2005). There are additionally the potential negative societal effects of animal consumption, particularly in terms of poor human attitudes toward animals leading to poor welfare outcomes overall (John and Sebo 2020).

With wild animals being probably the second highest source of animal suffering—higher in numbers but with more opportunities for positive welfare experiences to offset their suffering—investigating ways to manage wild animals to remove many of the negative experiences is another research priority. As mentioned, some advocate for the reduction in wild animal numbers as the best way to reduce suffering (Tomasik 2017), however, if it is possible instead to switch net-negative to net-positive lives, this will be a superior intervention than simply removing such lives. Currently, discussions of intervening on wild animal welfare are hampered by the sheer complexity of the task—we are famously terrible at making ecosystem changes without hosts of downstream negative effects. However, the more we know, the more possible it will be to do so, and perhaps aiming for a future in which we have the knowledge and ability to manage all wild animal populations for their maximal welfare would be ideal.

Another method would be to ensure that all the animals who do exist—captive or wild—are capable of increased wellbeing, not just through better life conditions but through use of technologies that make these animals capable of experiencing more pleasure (and/or less suffering). This could include selectively breeding or genetically engineering animals to

have a greater capacity for total pleasure, and/or an ability to take more pleasure in the conditions under which they usually find themselves. Some possible methods to achieve this include carefully managed gene-drives to introduce and spread genes that enhance welfare (Liedholm 2019), use of enhancement drugs that increase pleasure or take away suffering—discussed in the human case (Veit 2018; Veit et. al 2020) but so far given little attention for animals—or even development of technology for applying stimulation to the pleasure centres in the brain, which does not appear to be subject to diminishing marginal utility as other types of pleasures do (Ng 1997).

Related to this could also be engineering animals for reduced suffering. There is a small but growing literature regarding use of genetic engineering to create so-called ‘diminished’ animals who lack some of the species-typical capacities that currently create frustration and suffering, including the most extreme case ‘animal microencephalic lumps’ that completely lack sentience (Schultz-Bergin 2017). However, as this method removes the possibility of good lives and positive values, it will not end up creating the highest expected utility except in cases where the suffering would otherwise be inevitable. Where there are opportunity costs of directing resources away from creating or supporting otherwise happy animals, this would reduce total value.

These types of intervention may also interact with the changing animal numbers. Say, in the future, we are capable of creating very happy animals through use of chemical intervention or genetic engineering—this would then give us a reason to try to create and maintain as many animals as possible that are capable of experiencing this. One suggestion resulting from this is that this then might give us reason to try to maintain factory farms, as these are capable of holding the highest densities of animals, if the suffering currently experienced in such setups would be replaced with the types of pleasure described. However, this would only be true if we took factory farms as the best way of housing and keeping happy animals. This may be the case if we think that the economic incentives of using animal products would offset the costs of animal maintenance, but as we have discussed, it is also likely that there are many other setups and housing types that would, in actuality, be better for keeping large numbers of happy animals.

Some of these interventions may not be appealing to those with a less utilitarian approach to animal ethics, who have more of a concern for other values in animal lives, such as authenticity, or naturalness. We think there are reasons to question the role of such values: for instance, there is no obvious link between naturalness and moral value (Browning 2020)—pain, suffering, and extinction are, after all, perfectly natural phenomena and yet it is precisely these that we would most wish to avoid. It is important in any view of animal ethics to avoid anthropocentrism by considering what the animals themselves care about rather than imposing a human-centred view of what is valuable. Whereas humans might object to the prospects of being prescribed mood-enhancing drugs on grounds of authenticity, autonomy, or consent (see Veit 2018), there is less reason to think that such reasons would apply to animals. For example, as we have argued elsewhere, there is little reason to think that freedom must matter intrinsically to animal welfare (Browning and Veit 2020; 2021). Regardless, it is not our intention here to take a strong stance on views in animal ethics, and we take it to be the case that even where one is not in favour of some of these specific utilitarian actions, this still leaves a range of plausible interventions that undeniably improve animal lives on any account.

3.3 Value change

We have described a number of possible actions for improving the long-term future for animals, in terms of changing both the size and the sign of this future. However, more consideration is needed regarding which are likely to be the most effective actions for future benefit. When thinking about the long-term future, effective interventions will be those that persist for a long time, and are robust in the face of potential changes in future conditions. This can be framed in terms of attractor states—those states of the world that, once entered, are likely to continue for a long period (Greaves and MacAskill 2019). There are many potential attractor states, and some will be of more value than others. If our actions now can affect the probability that we enter a better rather than a worse attractor state, this will have ongoing effects. If we want to take the best possible actions for long-term future value, then focusing on the best attractor states is most likely to have the highest ongoing value. One potential action that we think may serve as an attractor state for the future of animals is change in human values and attitudes towards animals and their treatment, in terms of both individual and larger-scale institutional values.

Interventions to alter human attitudes and values could form a type of trajectory change from one attractor state to another, such that future policy and behaviour will be different. We can plausibly influence the direction of individual and institutional values toward those most likely to have positive future impacts for animals. This would include any action to ensure that future humans, particularly those with political power, hold attitudes that promote positive treatment of animals. Such changes will potentially have wide-ranging effects across all the domains we have described, ensuring their implementation and maintenance.

If our current stage in time is a particularly notable time in which we are about to see some form of ‘value lock-in’ (MacAskill 2020; 2022), where previously flexible or pluralist values give way to a single or rigid set of values that persist over a long timescale, then it is worth channelling resources now to ensure that these values are those that ensure good lives for future animals (and humans). This may be particularly likely if we think we are on the cusp of programming future superintelligent or autonomous AIs that will have a large influence on politics and society (Bostrom 2014). The values with which we program such AI systems can affect animals as much as they do humans, and the values they receive now may have a strong future influence on the conditions of animal lives. Though some scepticism has been expressed regarding how easily we may accurately represent animal interests (Ziesche 2021), even just ensuring that non-speciesist and ‘animal friendly’ attitudes are included should help ensure animals are given appropriate consideration.

One form such an attitude change could take would be a moral circle expansion—i.e. widening the circle of beings recognised as subjects of moral concern. This could include sentient animals and, potentially, other beings such as sentient AIs, which have recently received a surge of attention in the longtermist literature (John and Sebo 2020). Previous moral circle expansions based on shared humanity have provided increased protections for marginalised groups, and further expansion based on sentience would provide protection for animals, including farmed and wild animals (Anthis and Paez 2021). Moral circle expansion is compatible with a range of ethical frameworks; all it requires is that the type

of consideration that is currently offered only to humans would also be offered to other sentient animals. In practice this should mean that animal suffering counts for much more than it currently does; and many harmful practices could not continue. Depending on whether or not our current trajectory is already leading us to such an expansion, we could aim either to direct or to speed up expansion, both of which can have long-term benefits (Anthis and Paez 2021).

Even where one takes most of the future value to instead be realised by digital minds, it could be argued that this provides a reason to take animals seriously now, as neglecting animal interests may work to lock in anthropocentric values that would make it impossible in the future to ensure the consideration of the interests of other types of non-humans. It is possible that extending our moral circle to include other sentient animals may be a prerequisite for further extension towards sentient AI, the moral value of which is still not widely considered by either philosophers or the public. Such moral circle expansion, to include all sentient beings, would also require increased research into sentience and its neurological, cognitive, and functionalist basis to settle questions about which animals and AIs are sentient, and what their experiences are like. In light of the potential that our present is a particularly influential time where there is a higher risk of locking in harmful values into our institutions and legal frameworks, particularly when considering current work with AI, there may be a matter of urgency for pushing moral circle expansion right now, to ensure that it happens at all. Making sure that animal interests are included will help guard against cementing existing anthropocentric biases.

A concern about interventions based in human attitude change is whether or not they would actually work to improve the situation for animals. For instance, some take the historical evidence to speak *against* a correlation between value change and welfare—while we are arguably living in a time in which we hold the most animal-positive views, worldwide animal welfare is at its worst. However, this take misrepresents the current situation, and the relevant comparison class. The worsening state of animal welfare is largely due to increasing numbers being held in factory farms, which is a function of increasing population size and the spread of intensive farming techniques into new populations (Torrella 2021). If we hold fixed this increase in population size and look at the counterfactual situation regarding values—one in which societal attitudes towards animals had remained largely fixed, or gotten worse—then it is highly likely that the current situation would be even worse than it is now. It seems like we are closer than ever before to ending the practice of intensive farming, and human value change has been leading changes in practices, such as increasing adoption of bans on veal crates, sow stalls, and battery cages for chickens, that are an overall net benefit for welfare.

What is important is that we examine the causal links between societal values and the conditions of animal housing and husbandry. This will allow us to determine where best to target our interventions to create lasting and relevant value change. Whether individual changes of attitudes or wider structural change and improved institutional decision-making is more important may depend on what are the dominant mechanisms for value change, which is a matter for further research (see e.g. Sebo 2022). These should also not be taken as exclusive options, and indeed will often be complementary in that critical mass of individual lobbying or purchasing decisions, as well as research and policy advice, will influence institutional change.

4 Conclusion

The longtermist paradigm holds that the actions expected to produce the most good are those that have their effects in the long-term future. Here, we have argued that for the same reasons this argument is applied to considerations regarding humans, the wellbeing of future animals should also be given serious consideration when thinking about the long-term future. In fact, there is an interesting parallel between general longtermist thinking and an emphasis on the importance of animal welfare. Both are situations in which the group concerned (far-future populations or animal populations) compose a vast majority yet their interests are subsumed to the interests of a small majority (i.e. humans, or members of current and nearby future generations) and where the individuals concerned lack any political representation for their interests. There is thus a natural alignment between longtermism and more traditional animal advocacy, with promise for further collaboration. Studying the methods for creating successful change for consideration of the interests of animals who cannot make their voices heard may help us to likewise influence political institutions to take future humans, animals, and sentient AIs into account.

As well as including animals within longtermist thinking, we should additionally recognise the possibility that in some cases their aggregate interests may even dominate, due to their greater numbers and greater possible suffering. Even if one wants to resist this and maintain an anthropocentric priority, it is clear we should be giving animals much more consideration than is currently the case. We have not attempted to quantify the size of these effects or look at the relative calculations of the expected value contained in future animal versus human lives, but we suspect this could support an even stronger conclusion. That is, given the reasons we have presented, we actually have greater reasons to consider animals than humans: that our *best* actions in future will be those which benefit animals. Given the sheer numbers and level of suffering we see right now, it could even turn out to be the case that many short-term interventions to benefit animals could be more valuable than long-term interventions for human societies: for instance the lower bound estimate of a total of 10^{15} future people (Greaves and MacAskill 2019) is equal to the number of aquatic vertebrates existing *right now*. However, given the uncertainty surrounding the actual future numbers and level of suffering, as well as the comparative moral weight to humans, here we will content ourselves with the weaker claim that animals should at least be brought into deliberations regarding our best actions for the long-term future with a much greater weight than they are currently accorded. This is likely to significantly change the landscape of action prioritisation for the long-term future. We have described a range of potential interventions that change both the expected size and sign of the future, highlighting that actions targeted at changing human and societal attitudes are most likely to have a strong effect.

In some cases, the actions we have described (such as ending factory farming) will align with short-term priorities, but most often they are likely to focus on different initiatives—ending the most animal suffering *now* is not necessarily related to ending it in the long term, as is also true for human cases. Here, we sometimes have to push against our intuitions that we should be doing something now, if we accept the motivations for a longtermist worldview. Indeed, if we do not think we are living in a particularly influential period in time (i.e. one in which our interventions are likely to have unusually strong ongoing effects), then it may in fact be better to invest our resources such that they can be used for future interventions, rather than to take any direct action now (MacAskill 2020). Particularly where we are

currently uncertain about the specific changes that may end up being best for animals in the long run, empowering future generations to act on their behalf through shifting values and building up knowledge will be our current best action. However, some of the most important interventions discussed—such as institutional change and moral circle expansion, will have immediate as well as long-term effects.

The upshot of this chapter is not to advocate some specific action/s, but to call for the inclusion of animals in deliberations about the long-term future and which actions we should be prioritising for greatest gain. Importantly, it is a call for further research. While these questions are still uncertain, we should be gathering information such as the impact of attitude change on future behaviour, the net balance of pleasure/suffering in wild animals, and the likely future numbers of animals and humans. The mere assertion that we have little knowledge about how to improve the lives of animals is not enough to exclude them from a longtermist view, since it is precisely here that all future knowledge about animal welfare should be included. By bringing animals into our considerations, we can be surer that we will be making decisions that will have the best actual long-term impact, and it is our hope here that we have shown that the interests of animals should play a much larger role in longtermist thought and writing.

References

- Anomaly, J., Browning, H., Fleischman, D., and Veit, W. (2024), 'Flesh without blood: The public health benefits of lab-grown meat', *Journal of Bioethical Inquiry*, 21(1), 167–175.
- Anthis, J. R. and Paez, E. (2021), 'Moral Circle Expansion: A Promising Strategy to Impact the Far Future', in *Futures* 130: 102756.
- Bar-On, Y. M., Phillips, R., and Milo, R. (2018), 'The Biomass Distribution on Earth', in *Proceedings of the National Academy of Sciences* 115/25: 6506–6511.
- Beckstead, N. (2019), 'A Brief Argument for the Overwhelming Importance of Shaping the Far Future', in H. Greaves and T. Pummer (eds.), *Effective Altruism: Philosophical Issues* (Oxford University Press), 80–98.
- Bostrom, N. (2003), 'Astronomical Waste: The Opportunity Cost of Delayed Technological Development', in *Utilitas* 15/3: 308–314.
- Bostrom, N. (2013), 'Existential Risk Prevention as Global Priority', in *Global Policy* 4/1: 15–31.
- Bostrom, N. (2014), *Superintelligence: Paths, Dangers, Strategies* (Oxford University Press).
- Browning, H. (2020), 'The Natural Behavior Debate: Two Conceptions of Animal Welfare', in *Journal of Applied Animal Welfare Science* 23/3: 325–337.
- Browning, H. (2023), 'Welfare Comparisons Within and Across Species', in *Philosophical Studies* 180: 529–551.
- Browning, H. and Veit, W. (2020), 'Confined Freedom and Free Confinement: The Ethics of Captivity in *Life of Pi*', in Á. T. Bogár and R. S. Szigethy (eds.), *Critical Insights: Life of Pi* (Salem Press), 119–134.
- Browning, H. and Veit, W. (2021), 'Freedom and Animal Welfare', in *Animals* 11/4: 1148.
- Browning, H. and Veit, W. (2023), 'Positive Wild Animal Welfare', in *Biology and Philosophy* 38/14: 1–19.
- fishcount.org.uk. (2019), 'Fish Count Estimates', *fishcount.org.uk*, <http://fishcount.org.uk/fish-count-estimates-2> (accessed 25 February 2021).
- Franks, B., Ewell, C., and Jacquet, J. (2021), 'Animal Welfare Risks of Global Aquaculture', in *Science Advances* 7/14: eabg0677.
- Greaves, H. and MacAskill, W. (2019), 'The Case for Strong Longtermism', GPI Working Paper No. 7-2019 (Global Priorities Institute, Oxford University), https://static1.squarespace.com/static/5506078de4b02d88372eee4e/t/5f1704905c33720e61cd3214/1595344019788/The_Case_for_Strong_Longtermism.pdf (accessed 25 February 2021).
- Gruen, L. (2011), *Ethics and Animals: An Introduction* (Cambridge University Press).
- Harrison, R. (1964), *Animal Machines: The New Factory Farming Industry* (Vincent Stuart).
- Horta, O. (2010), 'Debunking the Idyllic View of Natural Processes: Population Dynamics and Suffering in the Wild', in *Télos* 17/1: 73–88.

- Iglesias, A. V. (2018), 'The Overwhelming Prevalence of Suffering in Nature', in *Revista de Bioética y Derecho* 42: 181–195.
- John, T. and MacAskill, W. (2021), 'Longtermist Institutional Reform', in N. Cargill and T. John (eds.), *The Long View* (FIRST), 44–60.
- John, T., and Sebo, J. (2020), 'Consequentialism and Nonhuman Animals', in D. W. Portmore (ed.), *Oxford Handbook of Consequentialism* (Oxford University Press), 564–591.
- Liedholm, S. E. (2019), *Persistence and Reversibility* (Wild Animal Initiative), https://static1.squarespace.com/static/5f04bd57a1c21d767782adb8/t/5f160c91bc0bff4abe964d5a/1595280529848/WAI_PersistenceAndReversibility_Dec2019.pdf (accessed 25 May 2021).
- MacAskill, W. (2020), 'Are We Living at the Hinge of History?', GPI Working Paper No. 12-2020 (Global Priorities Institute, Oxford University), https://globalprioritiesinstitute.org/wp-content/uploads/William-MacAskill_Are-we-living-at-the-hinge-of-history.pdf (accessed 15 March 2021).
- MacAskill, W. (2022), *What We Owe the Future* (Hachette).
- Matheny, G. and Chan, K. M. A. (2005), 'Human Diets and Animal Welfare: The Illogic of the Larder', in *Journal of Agricultural and Environmental Ethics* 18/6: 579–594.
- Mogensen, A. (n.d.), 'Staking Our Future: Deontic Longtermism and the Non-identity Problem', GPI Working Paper Series (Global Priorities Institute, Oxford University), https://globalprioritiesinstitute.org/wp-content/uploads/2019/Mogensen_Staking_Our_Future.pdf (accessed 15 March 2021).
- Ng, Y. K. (1995), 'Towards Welfare Biology: Evolutionary Economics of Animal Consciousness and Suffering', in *Biology and Philosophy* 10/3: 255–285.
- Ng, Y. K. (1997), 'A Case for Happiness, Cardinalism, and Interpersonal Comparability', in *The Economic Journal* 107/445: 1848–1858.
- Ord, T. (2020), *The Precipice: Existential Risk and the Future of Humanity* (Bloomsbury).
- Russell, A. (2023), 'How Many Vegans Are There in the World?', *WT Vox*, <https://wtvox.com/lifestyle/2019-the-world-of-vegan-but-how-many-vegans-are-in-the-world/> (accessed 30 November 2022).
- Schultz-Bergin, M. (2017), 'The Dignity of Diminished Animals: Species Norms and Engineering to Improve Welfare', in *Ethical Theory and Moral Practice* 20/4: 843–856.
- Sebo, J. (2022), *Saving Animals, Saving Ourselves* (Oxford University Press).
- Shriver, A. (2022), 'Why Neuron Counts Shouldn't Be Used as Proxies for Moral Weight', *EA Forum*, <https://rethinkpriorities.org/publications/why-neuron-counts-shouldnt-be-used-as-proxies-for-moral-weight> (accessed 30 November 2022).
- Šimčikas, S. (2020), 'Estimates of Global Captive Vertebrate Numbers', *Rethink Priorities*, <https://rethinkpriorities.org/research-area/estimates-of-global-captive-vertebrate-numbers/> (accessed 30 November 2022).
- Singer, P. (1975), *Animal Liberation* (Harper Collins).
- Tarsney, C. J. (2020), 'The Epistemic Challenge to Longtermism', working paper (PhilPapers), <https://philpapers.org/archive/TARTEC-2.pdf> (accessed 15 March 2021).
- Thorstad, D. and Mogensen, A. (2020), 'Heuristics for Clueless Agents: How to Get Away with Ignoring What Matters Most in Ordinary Decision-Making', GPI Working Paper No. 2-2020 (Global Priorities Institute, Oxford University), <https://globalprioritiesinstitute.org/wp-content/uploads/David-Thorstad-Andreas-Mogensen-Heuristics-for-clueless-agents.pdf> (accessed 15 March 2021).
- Tomasik, B. (2015), 'The Importance of Wild-Animal Suffering', in *Relations: Beyond Anthropocentrism* 3/2: 133.
- Tomasik, B. (2017), 'Habitat Loss, Not Preservation, Generally Reduces Wild-Animal Suffering', *Essays on Reducing Suffering*, <http://reducing-suffering.org/habitat-loss-not-preservation-generally-reduces-wild-animal-suffering/> (accessed 26 March 2021).
- Torrella, K. (2021), 'The Biggest Animal Welfare Success of the Past 6 Years, in One Chart', *Vox*, <https://www.vox.com/future-perfect/22331708/eggs-cages-chickens-hens-meat-poultry> (accessed 30 November 2022).
- Veit, W. (2018), 'Cognitive Enhancement and the Threat of Inequality', in *Journal of Cognitive Enhancement* 2: 404–410.
- Veit, W., Earp, B. D., Faber, N., Bostrom, N., Caouette, J., Mannino, A., Caviola, L., Sandberg, A., and Savulescu, J. (2020), 'Recognizing the Diversity of Cognitive Enhancements', in *AJOB Neuroscience* 11/4: 250–253.
- Višák, T. (2017), 'Preventing the Suffering of Free-Living Animals: Should Animal Advocates Begin the Killing?', in *Journal of Animal Ethics* 7/1: 78–95.
- Ziesche, S. (2021), 'AI Ethics and Value Alignment for Nonhuman Animals', in *Philosophies* 6/2: 31.