

## BETWEEN THE SPECIES

### Reducing Extreme Suffering for Non-Human Animals: Enhancement vs. Smaller Future Populations?

#### **ABSTRACT**

This paper argues that ethical views that place primary importance on the reduction of extreme suffering imply that, at least in theory, it can be better to allow enhanced non-human animals to come into existence rather than unenhanced non-human animals. Furthermore, they imply that it would be even better if no non-human animals came into existence at all. However, it is unclear, from the perspective of these ethical views, whether enhancement or reduction of future populations is the more effective strategy in practice, and whether it might even be better to instead pursue a seemingly more robust and less controversial third option of promoting greater concern for the suffering of non-human animals in the first place. In this paper, I seek to explore the different options from a practical perspective.

MAGNUS VINDING  
University of Copenhagen

Volume 23, Issue 1

Winter 2019

*<http://digitalcommons.calpoly.edu/bts/>*

## MAGNUS VINDING

In this paper I wish to examine what ethical views that place primary importance on the reduction of extreme suffering imply for the moral status of animal enhancement. In particular, I will argue that, according to such views, it can be better to allow certain enhanced non-human animals to come into existence rather than unenhanced non-human animals, yet the best would be if no non-human animals came into existence at all. I then proceed to discuss the implications of this conclusion for practical efforts to reduce extreme suffering for non-human animals.

**Preliminaries**

I would like to first clarify the scope and terminology of my argument. First of all, the argument I am making here pertains only to non-human animals. This is not to say that the argument cannot be made with respect to humans as well, but merely that this is not my concern in this essay. One reason for this is that non-human animals generally possess a very limited degree of moral agency. To a first approximation, most humans can be considered moral agents, whereas non-human animals cannot, at least not nearly to the same extent. And such moral agency does complicate discussions about the potential effects of enhancement and their ethical status, complications that can be ignored if we restrict our argument to non-human animals only. Note that I am not claiming that moral agents should be granted greater moral consideration than non-agents. The complications I seek to avoid by excluding moral agents mostly have to do with the ability that such agents have to help others.

Second, I should like to clarify what I mean by “ethical views that place primary importance on the reduction of extreme suffering”. By this I mean all ethical views according to which an action’s propensity to reduce extreme suffering

## MAGNUS VINDING

is the most important, yet not necessarily the only, criterion by which we evaluate its value and ethical status. This evaluative criterion comprises a consequentialist element of a moral theory. If this element stands alone, i.e. if an action's propensity to reduce extreme suffering is the only criterion by which we evaluate its value and normativity, one gets a particular kind of pure negative utilitarian view. Yet one can also include it as an element in a larger set of values. The crucial point, however, is that the ethical views whose implications I seek to examine here are those according to which this element is granted an overriding status. That is, views according to which the reduction of extreme suffering has overriding value, whether we also consider other things valuable or not (I shall abbreviate this class of ethical views by VPES, standing for "Values that grant Primacy to the reduction of Extreme Suffering"). For example, one can combine an overriding concern for the reduction of extreme suffering with consequentialist views according to which knowledge and happiness are also valuable. According to such a pluralist view, an outcome with more knowledge and more happiness would, other things being equal, be better than an outcome with less of these. Yet outcomes with more extreme suffering would always be worse than those with less, even if the outcomes with less suffering contain more knowledge or happiness. It is in this sense that the reduction of extreme suffering is granted an overriding status by VPES.

I shall not provide an elaborate defense of this view of the moral status of extreme suffering here (for that, I recommend consulting the sources cited below). A brief motivating thought experiment in its favor is to imagine two different planets, one of which can contain all the good things we can possibly imagine, while the other can contain only beings who experience extreme suffering, e.g., beings who experience being eaten or

## MAGNUS VINDING

skinned alive. The question, then, is whether we believe that adding good things, including happy lives, to the first planet can justify adding more extreme suffering to the other planet. VPES can be justified with reference to the intuition that no amount of good things added to the good planet can justify adding more extreme suffering to the bad one. Or, phrased more generally, that no amount of good things added to the world can justify adding more extreme suffering to it. This intuition can in turn be justified with reference to other intuitions about the moral importance of reducing extreme suffering versus promoting good things, such as happiness. For example, some (Popper 1945, chap. 9, note 2), have argued that there is a moral urgency – “a direct moral appeal for help” – in reducing suffering while there is no similar urgency or moral importance in increasing the happiness of the already happy. And even if one thinks other things do carry some moral urgency, such as increasing happiness and increasing knowledge, it seems plausible that this urgency is not comparable to the moral urgency of reducing extreme suffering. Karl Popper argued for this exact claim with respect to happiness: “. . . the promotion of happiness is in any case much less urgent than the rendering of help to those who suffer . . .” (Popper 1945, chap. 5, note 6). As for the idea that happy experiences or other good things can compensate for extreme suffering, one can question in what sense this could be the case. For example, some people claim to have experienced suffering so extreme that, according to them, no amount of good things could ever compensate for it (Vinding 2018a, chap. 7), which raises serious questions of how, and by whose standards, positive experiences or other positive goods are supposed to be able to compensate for these negative states. This is not obvious. Indeed, it is a widely held intuition that the most extreme states of suffering cannot be counterbalanced by positive experiences (Gloor & Mannino 2016, section II).

## MAGNUS VINDING

A fuller defense of the view that reducing extreme suffering has moral primacy can be found in (Vinding 2018a, chap. 6-7); (Vinding 2019, part I); (Leighton 2011, chap. 9); and (Gloor & Mannino 2016). It is worth noting that this view bears similarity to ethical views found within the tradition of Buddhism (Goodman 2009, 101), as well as to the view defended in Jamie Mayerfeld's *Suffering and Moral Responsibility* (Mayerfeld 1999, chap. 4-6). Indeed, a foremost moral priority to extreme suffering can be grounded in many different premises, such as by appeal to a prioritarian position or to the intrinsic badness of extreme suffering itself, and these premises may be combined or stand alone (Vinding, 2019, part I; Mayerfeld 1999, pp. 149-152). Furthermore, to underscore the relevance of examining the implications of VPES, it is worth noting that a recent survey in which 14,866 people were asked what goals they thought should be strived for by a future civilization, the goal that was favored more than any other, i.e. by a significant plurality, was minimizing suffering (Future of Life Institute 2017).

Third, I should also clarify that my argument is restricted to only concern beings who have not yet been brought into existence. There are three reasons why I have chosen this restriction: (1) because it allows us to steer clear of many complications that arise concerning the ethics of altering or killing existing beings versus altering or preventing the existence of future (merely) potential beings; (2) because future beings who have not yet been born likely comprise the vast majority of the beings we are able to impact with our actions (Beckstead 2013); and (3) because, even if we disregard the previous reason, our discussion concerning enhancement pertains primarily to future beings, since the extent to which we can change populations of currently living non-human animals, via enhancement or otherwise, is likely limited compared to

how much we will be able to change them in the future, when humanity will presumably be significantly wealthier and more technologically capable.

Lastly, I should clarify what I mean by “enhanced non-human animals” in this paper. By this I simply mean non-human animals who have been altered by humans (e.g., via gene editing or selective breeding) primarily for the purpose of benefiting these beings themselves (and hence enhancement, as I use the term, also encapsulates what is usually referred to as disenchantment, such as the reduction or removal of the ability to experience suffering). This definition excludes beings who have been altered purely for human benefit. It does not, however, exclude beings who do not in fact gain a net benefit, or who are even harmed, by the attempt to benefit them. In other words, the loaded term “enhanced non-human animals” neglects an important possibility which should not be overlooked, namely that human alterations of non-human animals, even when they are well-intentioned, can go wrong.

### **An Enhanced Population Can Be Better; No Population is Best**

There are two claims I wish to argue for here. The first is that, according to VPES, it can be better, at least in theory, to allow enhanced non-human animals to come into existence rather than unenhanced non-human animals. To realize this, we can imagine a scenario where a fixed number of non-human animals will come into existence in the future (on all of Earth, say). Provided that we can successfully alter this future population in such a way that it will contain less extreme suffering than it otherwise would—for example, by significantly down-regulating the pain-sensitivity of all individuals in the population (Pearce 2015, 157; Pearce 2016)—then such an al-

## MAGNUS VINDING

teration would, other things being equal, be normative according to VPES. Whether such enhancements are indeed feasible for humanity to realize for a large population of non-human individuals in practice, both technologically and sociologically, is another issue, one we shall return to below.

The second claim I wish to establish here is that, even in the best case, an enhanced future population of non-human animals would still, according to VPES, be strictly worse than no population of non-human animals at all, other things being equal. That is, between an idealized replacement option where unenhanced non-human animals are replaced, within one generation, by enhanced non-human animals, and a cessation option where no new beings are brought into existence, the latter is strictly better according to VPES.

This holds true since a population of zero non-human animals entails no extreme suffering, or indeed any suffering at all, for such beings, whereas this would not be true, even in the best possible case, of a future enhanced population. For even if it is possible to eventually get to a point where we have altered non-human animals so much that they can no longer experience extreme suffering, it is clearly not possible to get to that point in a single generation. For example, in any population of considerable size, accidents that cause severe bodily harm are bound to happen (e.g., severe burns or crushing of a body part), and it is not plausible that extreme suffering could be made impossible in the face of such accidents in just a single generation. Beyond that, there is also the problem of predation (McMahan 2016) and the extreme suffering it entails, which would also be impossible to phase out in just a single generation, or even a few.

## MAGNUS VINDING

Furthermore, even if we were to reach a stage where we have altered non-human animals so much that we are reasonably confident that they cannot experience extreme suffering, there would still be a risk that they in fact can and will, a risk which, even if quite small, would be much smaller still if there were no non-human animals at all. Indeed, this risk seems bound to remain significant, since non-human animals, unlike humans, cannot communicate their subjective states as well and in as much detail as humans can, both in terms of how they are feeling, and what the source of a given feeling seems to be. It seems doubtful, for example, that we would be able to assert beyond a reasonable doubt that a group of small fish is unable to experience extreme suffering merely by observing their behavior and knowing their full genetic profile. Thus, even in the best possible case where we have gathered as much information as we can, it seems that there would still be considerable uncertainty about the experiential states of enhanced non-human animals, and hence considerable risk of the realization of states of extreme suffering in these beings. (There would arguably be significantly less uncertainty in the case of humans, and this difference between human and non-human beings constitutes a reason, though by no means a decisive one, to favor a population of enhanced humans rather than of enhanced non-humans from the perspective of VPES; another such reason is the fact that non-human individuals and their suffering are generally granted far less moral concern by humanity than are human individuals, implying that non-humans face a greater risk of being harmed by humans). This risk should be minimized according to VPES, which is another reason why, even in the best case, and even disregarding the point made above concerning the impossibility of abolishing extreme suffering via enhancement in a few generations, a zero population would still, other things being equal, be strictly better than a population of en-



hanced non-human animals according to VPES. (It is worth noting that views other than VPES can favor a similar conclusion; for instance, (Ng 1995) and (Horta 2015) argue that suffering and disvalue prevail in the lives of most non-human animals, which, if true, would imply that a zero population would be better than the current one according to many different value systems, while (Vinding 2016) argues that merely applying a non-speciesist position on procreative ethics implies that it would be better if the vast majority of non-human animals were never born.)

In sum, a future population of enhanced non-human animals would, even in the best case, entail significant amounts, as well as risks, of extreme suffering, which, other things being equal, renders such a population worse than no population according to VPES. And since this conclusion also applies to any given sub-population within a larger population, it follows that a smaller population, of both enhanced and unenhanced non-human animals, is, other things being equal, better than a larger one.

### **Practical Implications: Two Controversial Options**

These conclusions are all rather theoretical, however, as they do not factor in any considerations about the practical feasibility of these two options of (1) enhancing future populations of non-human animals and (2) reducing or phasing them out altogether. For even if one accepts the evaluative conclusion we have drawn above at the theoretical level, one may still maintain that the best thing we can do to reduce extreme suffering, in practice, is to focus on enhancement rather than on phasing out or reducing populations of non-human animals.

## MAGNUS VINDING

This view is held by negative utilitarian David Pearce, who focuses primarily on the enhancement of future populations of non-human animals so as to gradually reduce, and eventually abolish, their suffering throughout the living world (Pearce 2007; 2015; 2016). He considers this the most promising and feasible way to reduce the suffering, including extreme suffering in particular, of non-human animals, all things considered. This is in part because he considers the option of enhancement, particularly via so-called CRISPR-based gene drives, a uniquely cheap way to reduce the suffering of non-human animals (Pearce 2016), but also in large part because he believes most people are more receptive to the idea of enhancing non-human animals than they are to the idea of phasing them out or significantly reducing their numbers (personal communication).

In contrast, Brian Tomasik, another negative utilitarian who has also written extensively about how to best reduce the suffering of non-human animals, believes the most promising approach to focus on is instead to reduce populations of non-human animals (Tomasik 2015c, 144; 2016a; 2016b). In particular, Tomasik views the reduction of natural habitats as a promising way to achieve such population reductions, which he considers more realistic than Pearce's strategy of encouraging humanity to altruistically pursue enhancements of non-human animals that have no benefit for humanity itself (Tomasik 2016a, 2016b). Indeed, Tomasik notes that he believes, with "99+% probability", that widespread altruistic efforts of the sort Pearce proposes to help wild animals will never come to fruition (Tomasik 2016b).

In the context of this paper, it is worth examining these respective proposals made by Pearce and Tomasik, including the reasoning behind them, in depth. For they are, to my knowl-

## MAGNUS VINDING

edge, the only concrete intervention proposals that have been defended with respect to the particular goal of minimizing the extreme suffering of future non-human animals. The underlying question here—which practical actions can be expected to best reduce extreme suffering?—is, it should be noted, an empirical question, one that depends on various factors which can be broken down and examined individually.

For one, there is the question of receptivity. Which idea would people be most willing to accept? As noted above, Pearce's position is in part animated by the belief that people will generally be more open to the idea of enhancement than they are to pure population reduction. Tomasik's position, on the other hand, rests more on the belief that most people would be too selfish and anthropocentric to support such altruistic efforts. And there are indeed reasons to be skeptical of the idea that enhancement would be more welcome than reductions of future populations. One such reason is that enhancement efforts of this kind seem bound to entail a significant amount of experimentation on non-human animals, which may be opposed by many. For although some studies indicate that a majority of people support such experimentation "so long as it is for [human] medical research purposes and there is no alternative" (Clemence & Leaman 2016, 4; Souza et al. 2017, 112), it should be noted that (1) the opposition against such experimentation is still significant, and (2) it is unclear whether there would be similar support for experimentation aimed toward reducing the suffering of non-human animals.

Beyond that, there is the fact many people still reject genetically modified (GMO) foods, despite many scientific studies concluding that they are as safe to consume as non-GMO foods (National Academies of Sciences 2016), which also constitutes

## MAGNUS VINDING

at least a weak reason to expect resistance to the idea of modifying non-human animals for their own sake. As for whether many people will ever be willing to benefit non-human individuals in the first place, there arguably is some evidence that people are in fact already supportive, at least in a generic sense. For example, a 1999 poll conducted in the United States found that 76 percent of people agreed that a “[non-human] animal’s right to live free of suffering is just as important as a [human] person’s right to live free of suffering” (Appleby 2008, 249).

Tomasik’s proposal, however, is in large part animated by the view that efforts to reduce or enhance future populations of non-human animals for altruistic reasons is unrealistic. Therefore, rather than trying to appeal to human altruism, Tomasik’s view is that the best thing to do is to push for interventions that reduce non-human populations, and that humans already want anyway for self-interested reasons, such as replacing grass lawns by gravel lawns or solar panels (Tomasik 2015a). As Tomasik writes:

Reducing wild-animal populations can be done by working with the grain of human selfishness rather than against it. For this reason, reducing populations seems plausibly more impactful as a way to prevent animal suffering in the short run than developing expensive, high-tech solutions that won’t arrive for many years and that humans have no selfish reason to deploy (Tomasik, 2016c).

Yet elsewhere, Tomasik argues that it nonetheless could make sense to devote some resources to enhancement approaches:

## MAGNUS VINDING

It would be a shame if people rested all their hopes on gene drives and didn't pursue other approaches to help wild animals in the short run. But exploring gene drives as one "high-risk stock" within a diversified portfolio of compassionate-biology research topics seems reasonable (Tomasik 2016b).

Thus, it need not be either-or, as one can clearly support efforts to both reduce and enhance future populations of non-human animals simultaneously, which Pearce does too, as he advocates for cross-species fertility regulation to prevent non-human animal populations from increasing (Pearce 2016). And so rather than being a matter of the one approach versus the other, the question is arguably just what level and form the emphasis on these two respective strategies should ideally take—a question that in turn depends on some of the open questions that have been raised above, namely: how receptive will humanity be to these respective ideas? And how altruistic can we expect humans to be toward non-human animals in the future? These are empirical questions that can, and probably should (from the perspective of VPES), be examined closer.

These questions about human attitudes are not the only relevant questions, however. Another important question is how technologically and economically feasible these proposed solutions are. Covering land with gravel or solar panels, for instance, is expensive, and covering all of Earth in this way will not be technologically nor economically feasible for the foreseeable future, if indeed ever. And the point Pearce would make is that the same is not true for the enhancement strategies he is proposing:

## MAGNUS VINDING

Back-of-an-envelope calculation suggests the financial cost of a happy non-human biosphere [by means of enhancement, e.g., via CRISPR gene drives] would currently be several hundred million dollars – plus annual maintenance costs of perhaps several million dollars per year (Pearce 2016).

In other words: all sentient beings on Earth could, according to Pearce, have their suffering alleviated and eventually abolished relatively cheaply. This is a staggering and highly disputable claim, of course. Yet it does nonetheless seem plausible that this enhancement solution could potentially, provided that it is indeed technically feasible, be much cheaper than the strategy of covering land, at least for very large areas of land. And if Pearce is right about the technological feasibility and low price of such an endeavor, this also has implications for the level of altruism required to make such a project happen. It could turn out to be very little.

Whether Pearce's estimates of both the technological and economic feasibility of such an effort of global enhancement of non-human animals are realistic is yet another empirical question that can be examined further. And along with the empirical questions listed above concerning human attitudes, this question is critical to explore in order to settle which of the two options discussed here – basically, enhancement via technology versus population reduction by means of habitat reduction – is best to spend one's marginal resources on in order to reduce extreme suffering.

### **A Third Option?**

Might there, however, be a third option that is even better than the two options examined above? At a first glance, if one

## MAGNUS VINDING

only looks at the level of direct interventions, there do indeed seem to be only two main options for reducing the suffering of future non-human animals – namely, to either reduce or change their populations significantly. Yet this does not imply that a direct focus on either of these two options is optimal in practice, when we factor in more indirect actions we can take.

At the level of advocacy, there seems to be a real risk that advocating for enhancement and/or reduction of future populations of non-human animals could backfire and end up being counterproductive due to the controversial nature of these proposals. Indeed, it seems likely that many people who would otherwise support the underlying value that animates these proposals—i.e. the reduction of extreme suffering experienced by non-human animals—will nonetheless reject these controversial proposals for intervention, a rejection that seems even more likely if this underlying value has not been explained and argued for carefully in the first place.

This, then, tentatively suggests another approach to prioritize: to promote concern for the suffering, especially the extreme suffering, experienced by non-human animals, and to argue for the moral primacy of its reduction. Beyond the fact that people will likely be more receptive to such a message, a lack of concern for non-human suffering may also constitute the main bottleneck with respect to the realization of concrete interventions that prevent extreme suffering, such as those discussed above. Pearce, for instance, seems to view humanity's level of concern for the suffering of non-human animals, not technological or economic limitations, as the main bottleneck by far that prevents the realization of the enhancement scenario he argues for (Pearce 2007; 2016). Tomasik, too, seems to view such a focus on values as being uniquely promising, among

## MAGNUS VINDING

other reasons because: (1) it is more robust in the face of new information—promoting a set of values seems more likely to be positive for the realization of these values than promoting particular interventions that may turn out to be negative in light of new information; (2) one can build a movement around these values, which both seems more feasible than does building a movement centered around, say, habitat reduction, as well as more likely to be positive and flexible with respect to optimizing that set of values than a more specific focus on a particular intervention (e.g., one can thereby get more people to join the endeavor of figuring out which concrete interventions best help alleviate the suffering of non-human animals, and how to best reduce extreme suffering in general); and finally, (3) a focus on promoting VPES seems less replaceable than most other efforts one could pursue to reduce extreme suffering:

Most object-level projects that one might undertake are already being done by lots of people with various ideological positions. In contrast, there are very few people promoting suffering-focused ethical viewpoints. So efforts to promote suffering-focused ethics may have more counterfactual impact than promoting a more mainstream cause that's less specific to your values (Tomasik 2015b).

And one may argue that habitat reduction, although it is all but never considered an altruistic cause, indeed is an endeavor that is carried out by many more people than is advocacy for the moral significance of the extreme suffering of non-human animals, and hence that one may expect, by the argument in Tomasik's quote above, that using one's marginal resources to promote habitat reduction would be less effective than the promotion of concern for such suffering. Also because, even if



## MAGNUS VINDING

one believes that the best one can do, in terms of direct interventions, is to push for slightly more habitat reduction on the margin, it could be that focusing primarily on the promotion of concern for the suffering of non-human animals is the best way to gain support for such efforts in the first place.

The fact that it is such a clear point of agreement that the promotion of concern for the extreme suffering of non-human animals is beneficial with respect to VPES, and that it indeed may be what best addresses the most crucial bottleneck with respect to the alleviation of such extreme suffering, also comprises a reason to focus one's practical efforts most strongly on the promotion of such concern. In contrast, the utility of any given direct intervention (e.g. habitat reduction or enhancement via genetic engineering) is much less agreed upon, and also in part for that reason more controversial, which constitutes a reason to be more careful and hesitant about focusing mostly on such interventions (both due to considerations concerning cooperation between disagreeing agents, as well as considerations about epistemic humility in light of empirical disagreements, in this case about what best reduces extreme suffering in practice).

It should be noted in this context, however, that, from the perspective of VPES, it is not enough to merely promote increased moral consideration for non-human animals per se, since such increased consideration for these individuals, if not coupled with a strong concern for their suffering, may in fact lead to an increase rather than a reduction of suffering for these individuals. This is because other value systems may favor increasing the numbers of those beings who fall within the notional circle of moral consideration, even if such an increase happens at the cost of an increase in extreme suffering. In this way, generic

moral circle expansion can actually end up being a very bad thing from the perspective of VPES, and hence should likely, from this perspective, be supplemented strongly by arguments for the moral primacy of reducing extreme suffering (Vinding 2018a, chap. 9; Vinding 2018b; Tomasik 2015b).

Conversely, it is also true that merely promoting concern for extreme suffering is not sufficient either, as there is a risk that such increased concern will fail to pertain properly to the suffering of non-human animals, as opposed to relating almost purely to human suffering. Thus, from the perspective of VPES, both widening the moral circle and deepening concern for extreme suffering are necessary yet insufficient on their own (Vinding 2018a, chap. 9). How to best effect such changes in society in more concrete terms stands as an open question, yet producing and publicizing written pieces, lectures, and documentaries that argue for such a broader and deeper concern for extreme suffering seem promising options.

### **Conclusion**

In this paper I have argued that, even in the best case, a future population of enhanced non-human animals would entail significant amounts, as well as risks, of extreme suffering. This renders such a population worse than no population according to VPES, and a smaller population better than a larger one, other things being equal.

These theoretical conclusions concerning the implications of VPES are important and worth being clear about, not least because they imply that it would be better, according to VPES, to avoid increasing future populations (for instance, by spreading non-human animal life into space), even if we believed the risk that these populations will entail extreme suffering to be

## MAGNUS VINDING

very low. Nonetheless, these theoretical conclusions still do not clearly show which actions that are optimal to take in practice in order to best reduce the extreme suffering of future non-human animals, which stands as an open empirical question. The utility of focusing on concrete interventions directed at enhancing and/or reducing future non-human animal populations is unclear and dependent on various factors, some of which can be investigated empirically, such as people's attitudes toward such proposals, the level of altruism we can expect humans to exercise toward non-human animals, and the technological and economic feasibility of the concrete interventions in question. Such investigations may help inform which specific interventions are most worth focusing on, if any.

Finally, I presented some reasons why the best way to reduce extreme suffering for future non-human animals, in practice, may be to focus mostly on the strategy of promoting concern for such suffering in the first place, such as by arguing that its reduction has moral primacy. This seems a more robust strategy that one can more easily build a movement around compared to if one were to focus more purely on direct and controversial interventions, such as gene editing or habitat reduction. Beyond that, such increased concern may also be what best addresses the main bottleneck that prevents us from knowing and doing more at the level of direct interventions, and hence this indirect strategy could well be the best way to foster progress at this more direct level as well. In other words, rather than focusing mostly on enhancing or otherwise altering non-human animal populations, the best way to reduce the extreme suffering of non-human animals in the future may, at this point at least, be to focus most of our resources on enhancing humanity's concern for such suffering first and foremost, and to argue for the moral primacy of its reduction.

## References

- Appleby, Michael. 2008. Long distance transport and welfare of farm animals. Wallingford, UK Cambridge, MA: CABI.
- Beckstead, Nick. 2013. On the overwhelming importance of shaping the far future, PhD Dissertation. New Brunswick: Rutgers University.
- Clemence, Michael & Leaman, John. 2016. "Public attitudes to animal research in 2016". Ipsos MORI Social Research Institute. <https://www.ipsos.com/sites/default/files/publication/1970-01/sri-public-attitudes-to-animal-research-2016.pdf> (accessed September 18, 2018).
- Future of Life Institute. 2017. "The Future of AI – What Do You Think?" Futureoflife.org. <https://futureoflife.org/superintelligence-survey/> (accessed July 25, 2018).
- Gloor, Lukas & Mannino, Adriano. 2016. "The Case for Suffering-Focused Ethics". Foundational-research.org. <https://foundational-research.org/the-case-for-suffering-focused-ethics/> (accessed September 15, 2018).
- Goodman, Charles. 2009. Consequences of Compassion: An Interpretation and Defense of Buddhist Ethics. Oxford: Oxford University Press.
- Horta, Oscar. 2015. "The Problem of Evil in Nature: Evolutionary Bases of the Prevalence of Disvalue". *Relations: Beyond Anthropocentrism*, 3, no. 1: 17-32.
- Leighton, Jonathan. 2011. The Battle for Compassion: Ethics in an Apathetic Universe. New York: Algora Pub.

## MAGNUS VINDING

- Mayerfeld, Jamie. 1999. *Suffering and Moral Responsibility*. New York Oxford: Oxford University Press.
- McMahan, Jeff (2016) "The moral problem of predation", in Chignell, Andrew, Terence Cuneo, & Matthew C. Halteman. *Philosophy Comes to Dinner: Arguments on the Ethics of Eating*. New York: Routledge-Taylor & Francis.
- National Academies of Sciences, Engineering, and Medicine. 2016. *Genetically Engineered Crops: Experiences and Prospects*. Washington, DC: The National Academies Press.
- Ng, Yew-Kwang. 1995. "Towards Welfare Biology: Evolutionary Economics of Animal Consciousness and Suffering". *Biology and Philosophy*, 10, no. 3: 255-85.
- Pearce, David. 2007. "The Abolitionist Project". *Abolitionist.com*. <https://www.abolitionist.com/> (accessed July 11, 2018).
- Pearce, David. 2015. "A Welfare State for Elephants? A Case Study of Compassionate Stewardship". *Relations: Beyond Anthropocentrism*, 3, no. 2: 153-164.
- Pearce, David. 2016. "Compassionate Biology: How CRISPR-based 'gene drives' could cheaply, rapidly and sustainably reduce suffering throughout the living world". *Gene-drives.com*. <https://www.gene-drives.com/> (accessed July 11, 2018).
- Popper, Karl. (1945) 2011. *The Open Society and Its Enemies*. London: Routledge.

## MAGNUS VINDING

- Souza, Ana Paula Oliveira, et al. 2017. "Brazilian attitudes towards the use of animals in research (Atitudes de brasileiros no uso de animais em pesquisa)". *Braz. J. Vet. Res. Anim. Sci.*, 54, no. 2: 109-116.
- Tomasik, Brian. 2015a. "Convert Grass Lawns to Gravel to Reduce Insect Suffering". *Reducing-suffering.org*. <http://reducing-suffering.org/convert-grass-lawns-to-gravel-to-reduce-insect-suffering/> (accessed July 11, 2018).
- Tomasik, Brian. 2015b. "Reasons to Promote Suffering-Focused Ethics". *Reducing-suffering.org*. <http://reducing-suffering.org/the-case-for-promoting-suffering-focused-ethics/> (accessed July 20, 2018).
- Tomasik, Brian. 2015c. "The Importance of Wild-Animal Suffering". *Relations: Beyond Anthropocentrism*, 3, no. 2: 133-152.
- Tomasik, Brian. 2016a. "Habitat Loss, Not Preservation, Generally Reduces Wild-Animal Suffering". *Reducing-suffering.org*. <http://reducing-suffering.org/habitat-loss-not-preservation-generally-reduces-wild-animal-suffering/> (accessed July 11, 2018).
- Tomasik, Brian. 2016b. "Will Gene Drives Reduce Wild-Animal Suffering?" *Reducing-suffering.org*. <http://reducing-suffering.org/will-gene-drives-reduce-wild-animal-suffering/> (accessed July 11, 2018).
- Tomasik, Brian. 2016c. "Why I Don't Focus on the Hedonistic Imperative". *Reducing-suffering.org*. <http://reducing-suffering.org/dont-focus-hedonistic-imperative/> (accessed July 11, 2018).

## MAGNUS VINDING

Vinding, Magnus. 2016. "The Speciesism of Leaving Nature Alone and the Theoretical Case for 'Wildlife Anti-Natalism'". *Apeiron*, 8: 169-183.

Vinding, Magnus. 2018a. *Effective Altruism: How Can We Best Help Others?* North Carolina: The Neuroethics Foundation.

Vinding, Magnus, 2018b. "Moral Circle Expansion Might Increase Future Suffering". *Magnusvinding.com*. <https://magnusvinding.com/2018/09/04/moral-circle-expansion-might-increase-future-suffering/> (accessed October 9, 2018).

Vinding, Magnus, 2019 (forthcoming). *Suffering-Focused Ethics*. North Carolina: The Neuroethics Foundation.

MAGNUS VINDING



MAGNUS VINDING