

The Comparative Badness for Animals of Suffering and Death

Jeff McMahan
November 2014

1 Humane Omnivorism

An increasingly common view among morally reflective people is that, whereas factory farming is objectionable because of the suffering it inflicts on animals, it is permissible to eat animals if they are reared humanely and killed with little or no pain or terror. I will refer to methods of rearing and killing animals that are designed to avoid causing significant suffering or terror as *humane rearing*, and will use the label *Humane Omnivorism* to refer to the view that, while it is wrong to buy and eat meat produced by factory farming, it is permissible to eat meat produced through humane rearing.

Although humane rearing, when practiced scrupulously, does not cause animals to suffer, it does involve killing them quite early in their lives. Beef cattle have a natural life span of about 30-35 years but are normally killed at about three years of age. Pigs can live about 15 years but tend to be killed at about six months, while chickens can live about eight years but are killed less than a year after birth. The reason these animals are killed when young is that it is economically wasteful to invest resources in keeping them alive after they have reached their full size.

Because humane rearing involves depriving animals of most of the life they might otherwise have, Humane Omnivorism seems to presuppose that depriving animals of good experiences is not morally objectionable in the way that causing them to suffer is. It seems to assume that although their suffering matters, their lives matter much less, perhaps not at all.

Even those contemporary philosophers who are critical of practices that use animals for human benefit tend to focus their criticisms on the causation of suffering. They frequently quote with approval Bentham's well-known claim that "the question is not, Can they *reason?* nor, Can they *talk?* but, Can they *suffer?*" Less often quoted is Bentham's claim, in the same footnote, that if the practice of eating animals did not cause suffering, it would be morally unproblematic, for of the animals that are killed he says that "they are never the worse for being dead."¹ I assume that he would not similarly conclude that, if people who have been killed are also not the worse for being dead, killing people for pleasure is permissible as well provided there are no bad effects on the survivors. Presumably Bentham thinks that whatever the fundamental reason is for not killing persons does not apply to the killing of animals.

¹ Jeremy Bentham, *An Introduction to the Principles of Morals and Legislation* (Oxford: Clarendon Press, 1879). p. 311.

The doctrine of Humane Omnivorism is nicely manifest in the work of Temple Grandin, the well-known activist on behalf of animal welfare and the rights of the autistic. Grandin, herself autistic, apparently has a profound empathetic understanding of the emotions of certain animals, particularly cows, and has made great efforts to reform the practices of the meat industry so that animals in slaughterhouses can go peacefully and without fear to their deaths. Yet she apparently has no qualms about depriving animals of good experiences by killing them. "Using animals for food is an ethical thing to do," she says, "but we've got to do it right. We've got to give those animals a decent life, and we've got to give them a painless death. We owe the animal respect."² Yet the suggestion that killing an animal for pleasure when its life would otherwise be worth living is compatible with respecting it obviously requires some explanation.

Humane Omnivorism grants that the suffering of animals matters enough to make factory farming, and supporting factory farming by eating the meat it produces, wrong. But it does not have to concede that the suffering of an animal matters as much as the equivalent suffering of a person. It might instead claim that the suffering of an animal matters less, perhaps on the ground that the suffering of beings with lower moral status matters less than the equivalent suffering of beings with higher moral status. Consider, for the sake of illustration, a version of Humane Omnivorism based on the artificially precise assumption that the suffering of an animal matters a tenth as much as the equivalent suffering of a person. Like all versions of Humane Omnivorism, this version accepts that it is permissible to kill an animal as a means of enabling people to have the greater pleasure of a certain number of meals that include meat from the animal rather than the lesser pleasure of an equivalent number of meals without meat. This is true even when, if the animal had not been killed, it could have had many more years of life without significant suffering.

Suppose that one animal would provide meat for twenty meals and that twenty people would each get ten more units of pleasure from eating a meal with meat from the animal than they would have got from otherwise similar meals without the meat, though with some substitute plant-based food of equivalent cost and nutritional benefit. We might next ask how much suffering it might be permissible to cause a *person* to experience as a means of enabling twenty other people to experience ten units of pleasure each. The answer is surely that it would be permissible to cause a person at most only a tiny amount of suffering for this reason. Given a version of Humane Omnivorism that accords the suffering of an animal a tenth the weight of the equivalent suffering of a person, it follows that it would be permissible to cause an animal no more than ten times this tiny amount of suffering as a means of providing the ten units of pleasure to each of ten people. Ten times a tiny amount of suffering is a small amount of suffering. So this version of Humane Omnivorism implies that while it would be permissible to deprive an animal of many years of life as a means of providing

² http://en.wikipedia.org/wiki/Temple_Grandin. Last accessed 9 September 2014.

each of twenty people a certain amount of pleasure, it would not be permissible to cause that animal more than a small amount of suffering as a means of providing the same pleasure to the same people. According to this version of Humane Omnivorism, therefore, it is worse to cause an animal to experience a small amount of suffering than it is to kill the animal, even when killing it would deprive it of many years of life without significant suffering, so that the good life it would lose would be good for it by much more than the suffering would be bad for it. Many other possible versions of Humane Omnivorism make similar assumptions about the relative badness of suffering and death for animals.³ We can group these assumptions under the label *Suffering is Worse*. My aim in this article is to examine the claim that Suffering is Worse.⁴

I will begin, in section 2, by considering whether there is a general moral asymmetry between suffering and happiness. In section 3, I sketch an account of the misfortune of death that helps to explain why the loss through death of a certain amount of good life is generally a lesser misfortune for an animal than for a person. This account, therefore, provides some support for Suffering is Worse. But section 4 shows that the account is vulnerable to counterexamples and section 5 shows that one appealing way of avoiding the counterexamples threatens to deprive the account of its distinctive virtues. In sections 6 through 9 I explain how the account's apparently implausible implications can be avoided or at least mitigated by conjoining it with either of two plausible views about the morality of causing individuals to exist. Section 10 then considers the implications of the two conjoined views for the morality of abortion and section 11 summarizes their implications for Suffering is Worse and Humane Omnivorism.

2 Possible Defenses of the Claim that Suffering is Worse

Many people believe that there is a general moral asymmetry between suffering and happiness for persons as well as for animals – that is, that the reason not to cause or allow individuals to suffer is normally stronger than the reason to cause or allow individuals to experience a corresponding degree of happiness. Some people for example, accept a strong form of this view according to which pure benefits cannot on their own morally offset harms. These people think that while it can be permissible to cause a person to suffer, even without her consent, if that is necessary to prevent her from experiencing even greater suffering, it is not permissible to cause her to suffer without her consent to provide her with a benefit that would be good for her by more than

³ An alternative basis for Humane Omnivorism is an argument that Peter Singer calls the “replaceability argument.” See his *Practical Ethics*, 3rd ed. (Cambridge: Cambridge University Press, 2011), pp. 104-19. I will not discuss that argument here, though I have done so in “Eating Animals the Nice Way,” *Daedalus* (2008): 66-76.

⁴ I have written material relevant to this issue before, first in *The Ethics of Killing: Problems at the Margins of Life* (New York: Oxford University Press, 2002), pp. 182, 195-98, 199-203, 229-30, 474-75, and 487-93.

the suffering would be bad.⁵ The claim that Suffering is Worse in animals might simply be a corollary of this general asymmetry between causing suffering and bestowing benefits. For although we think of death as a harm, it does not involve suffering or anything intrinsically bad. It involves only the absence of pure benefits – the good experiences and activities that continued existence would have made possible.

This asymmetry is, however, difficult to reconcile with the fact that we all accept, and believe that it is rational to accept, trade offs between suffering and happiness. People are often willing to undergo considerable suffering to achieve benefits for themselves, even when they would not suffer from the absence of those good things. If challenged, they may say that the benefits outweigh the suffering – thereby implying that the suffering and the benefits can be roughly measured on a common scale and that the benefits would be good for them by more than the suffering would be bad for them. Some people think such judgments are subjective, in the sense of being objectively ungrounded preferences. But it seems objectively true that, for example, it is better for a person to experience a moment's mild suffering if that is necessary for her to have a long period of great happiness, even if she strongly prefers not to experience the suffering. It therefore seems a mistake to suppose that, when it is not possible to get a person's consent, it is permissible to cause him to suffer when that is necessary to prevent him from experiencing a greater harm, but not to enable him to have a greater benefit. I suspect that the explanation of our attraction to this view is that, while it is often uncontroversial that one instance of suffering is worse than another, so that one can often be confident that a harm one inflicts will be less bad than a harm one thereby prevents, there is greater diversity in what benefits people. In consequence, there is often considerable uncertainty about whether or to what extent a person would benefit from something from which some others would benefit, and thus about whether the presumed source of benefit would in fact compensate the person for the suffering that is a necessary means of bringing it about.

Another reason for skepticism about the moral asymmetry between causing suffering and providing benefits is that, although we tend to accept it in a range of cases, we reject its application to persons in exactly the kind of choice to which defenders of Humane Omnivorism think it applies in the case of animals – namely, choices between suffering and death. Again, if one saves a person's life, one does not prevent anything that is intrinsically bad for her; one merely enables her to have the benefits of continued existence. Yet in instances in which one can save a person's life only by causing her to suffer, we believe that it is permissible, and perhaps required, to save her, even when it is not possible to get her consent, provided that the net benefit to her in remaining alive would outweigh the suffering it would be necessary to inflict.

⁵ See, for example, Seana Shiffrin, "Wrongful Life, Procreative Responsibility, and the Significance of Harm," *Legal Theory* 5 (1999): 117-48.

This challenges the view that there is a general asymmetry between causing or allowing individuals to suffer and causing or allowing them to have pure benefits. If it is permissible to cause people to suffer, even when one cannot get their consent, to enable them to continue to have benefits rather than to die, it ought to be permissible as well to cause them to suffer to enable them to have more benefits rather than continuing to live with fewer, provided the increase in benefits would outweigh the suffering.

Some defenders of Humane Omnivorism seem to think it unnecessary to give arguments to show that the painless killing of an animal is justified by the benefits that people derive from eating it. They seem to assume, instead, that the painless killing of an animal that has been humanely reared requires no moral justification at all. But presumably that could be true only if the well-being or happiness of animals does not matter at all. Perhaps the best defense of that claim is the austere perfectionist view that the only pleasures that are accessible to most animals, such as eating, playing, and lying in the sun, are too “low” to matter. But while it may well be true that some higher forms of pleasure have lexical priority over certain lower ones, it is implausible to suppose that these lower pleasures have no value at all. If the lower pleasures did not matter, all that would matter in the well-being of infants and small children would be the avoidance of suffering, except insofar as their experience of the pleasures of which they are capable would be instrumental to their ability to experience higher pleasures later in life.

Someone might argue that because animals have a lower moral status, the lower pleasures do not matter in their lives, though they do matter in the lives of persons. That is, the moral status of animals is such that their lives do not matter, even though their suffering does. This, however, seems entirely ad hoc. It is hard to see what kind of rationale could be given for it. If it is compatible with their lower moral status that their suffering matters, it seems that their happiness should matter as well, at least to some extent. If, moreover, the suffering of animals mattered but their happiness did not matter in any way, there would be a moral presumption against causing or allowing any animal to exist; for even if it is not inevitable that any animal will experience suffering, it is inevitable that any animal will be at risk of suffering. On this view, therefore, if it is morally justifiable to cause or allow an animal to exist, and thus to be exposed to the risk of suffering, that must be because of the benefits to people of its existing or the costs to people of preventing it from existing. Even more implausibly, the view also seems to imply that, unless animals have a right not to be killed, which seems incompatible with a conception of their moral status that denies that their lives matter, there must be a moral presumption against allowing any animal to continue to exist. No matter how much happiness an animal’s future life would contain, if it would also contain some suffering, or if the animal would be at risk of suffering, it could be permissible to allow it to continue to exist only because of the benefits to people of its existence or the cost to people of killing it painlessly. But that seems false.

3 The Time-Relative Interest Account

The claim that Suffering is Worse is clearly not based on the idea that there is something especially bad about the suffering of animals. It is based, rather, on the view that death is *not* especially bad for animals. Part of the explanation of this is obvious: the life that an animal loses in dying is almost always inferior in quality and quantity to that of which a person is deprived by death. Given the orthodox and in my view correct view that death is bad for an individual primarily because it deprives her of further life that would have been good for her, it is natural to suppose that how great a misfortune an individual suffers in dying varies with the quality and quantity of the life she would otherwise have had. Because animals usually lose less, their deaths are usually less bad.

Many philosophers have argued that death is a misfortune for an individual in direct proportion to the amount of good life it prevents the individual from having.⁶ But this implies that the worst death that an individual can suffer is immediately after beginning to exist. If, for example, we begin to exist at conception, as most people seem to believe, the death of a zygote immediately after conception is the worst death anyone can suffer. Yet most people believe that if a sperm and egg were about to fuse in conception but were destroyed just prior to fusing, there would be no significant loss at all. But if that is right, it is hard to believe that if the zygote were to die immediately after the sperm and egg have fused, without experiencing even a flicker of consciousness, it would be the victim of a terrible misfortune: the worst death an individual could suffer. On any plausible view about when we begin to exist, it is intuitively implausible that death immediately after that is a greater misfortune for the individual at the time than, say, the death of a person at age twenty.

I have argued elsewhere that the extent to which death is a misfortune for an individual is a function primarily of two independent factors: (1) the amount of good life of which the individual is deprived by death and (2) the extent to which the individual at the time of death would have been psychologically connected to himself at those times in the future when the good things in his life would have occurred.⁷ On any plausible view of when we begin to exist, we come into existence either without the capacity for consciousness or in a psychologically rudimentary condition. At best we are then only very weakly psychologically related to ourselves in the future. We are unaware of having a future in prospect and thus have no future-directed desires or intentions; nor will we later have any memory of our present experience. There would be scarcely any difference of significance between our dying at that point and our never

⁶ See, for example, Fred Feldman, *Confrontations With the Reaper* (New York: Oxford University Press, 1994); Ben Bradley, *Well-Being and Death* (New York: Oxford University Press, 2011); and John Broome, "The Badness of Death and the Goodness of Life," in Ben Bradley, Fred Feldman, and Jens Johansson, eds., *The Oxford Handbook of the Philosophy of Death* (New York: Oxford University Press, 2013): 218-33.

⁷ For a detailed discussion and defense of the claims in the previous paragraph, this paragraph, and the following two paragraphs, see Jeff McMahan, *The Ethics of Killing: Problems at the Margins of Life* (New York: Oxford University Press, 2002), pp. 165-74.

having come into existence at all. But as we mature psychologically, we gradually become both more substantial as possible subjects of misfortune and more closely psychologically connected to ourselves as we will later be if death does not intervene. Until we reach a certain level of psychological capacity, death becomes a greater misfortune as we develop even though the amount of good life we have in prospect is steadily diminishing.

Like ourselves in the earliest moments of our lives, most animals are, throughout their lives, largely psychologically unconnected to themselves in the future. They live mostly in the present. So not only is the life they lose through death inferior in quality and quantity but they are also only weakly related to their possible future life in the ways that matter. The magnitude of the misfortune they suffer in dying is diminished accordingly. The strength of an animal's interest in continuing to live is, one might say, discounted for psychological unconnectedness between itself at the time of death and itself at the times at which it would have had good experiences in the future. I call this the *Time-Relative Interest Account* of the misfortune of death (TRIA for brevity). According to this account, the strength of an individual's present interest in some possible event reflects the degree to which it is rational to care for the individual's own sake now whether that event will occur. It does not necessarily reflect the way the event would affect the value of the individual's life as a whole.

The badness of causing an animal to suffer now is, by contrast, unaffected by diachronic psychological unconnectedness. Even if the suffering of an animal matters less because of its lower moral status, it may be worse for the animal to suffer now than to die now because the extent to which its losses through death matter is steeply discounted for psychological unconnectedness. This point can be illustrated with an example.

Suffering Now An animal has a condition that will soon kill it painlessly. One can save it but only in a way that will cause it moderate suffering beginning shortly and continuing for a few days. It will then live for five years in a hedonically neutral state, followed by ten years of comfort, during which it will experience some of the highest forms of happiness of which it is capable.

According to the TRIA, the animal's present interest in avoiding the immediate suffering is strong while its interest in experiencing greater happiness in the distant future is weak because it would be only very weakly psychologically connected to itself during that later time. Depending on the relative weights of the two factors (amount of happiness and degree of psychological unity), the TRIA could imply that it would be better to allow the animal to die. In this case, therefore, the TRIA supports Suffering is Worse.

4 Objections to the Time-Relative Interest Account

The TRIA is not, however, consistent in its support for Suffering is Worse. This is evident in a parallel example in which the temporal ordering of suffering and happiness is reversed.

Suffering Later An animal has a condition that will soon kill it painlessly. One can save it only in a way that will enable it to experience moderate happiness beginning shortly and continuing for a few days, after which it will live five years in a hedonically neutral state, followed by a few weeks of intense suffering before it finally dies. If one saves it now, there will be no opportunity to prevent it from suffering later.

If the TRIA can imply that in *Suffering Now* it is better to allow the animal to die, even though its future happiness would outweigh its immediate suffering, it can also imply that in *Suffering Later* it would be better to save the animal, even though its future suffering would outweigh its immediate happiness.

The TRIA's implication in *Suffering Later* seems quite counterintuitive.⁸ Many people may find its implication in *Suffering Now* implausible as well, though those committed to *Suffering is Worse* should find it plausible.

Elizabeth Harman finds the TRIA's implications in both cases implausible. She argues that the TRIA is vulnerable to the following counterexample.

Tommy is a horse with a serious illness. If the illness is not treated now and is allowed to run its course, Tommy will live an ordinary discomfort-free life for five years, but then Tommy will suffer horribly for several months and then die. If the illness is treated now, then Tommy will undergo surgery under anesthetic tomorrow. Tommy will suffer over the following two weeks, but not nearly as severely as he would five years from now. Tommy will be completely cured and will be able to live a healthy normal life for another fifteen years.⁹

Harman believes that it would be better for the horse to be treated. It is better for it to have suffering now together with greater happiness in the future (which conflicts with the TRIA's judgement in *Suffering Now*) rather than happiness now at the cost of greater suffering in the future (which conflicts with the TRIA's judgement in *Suffering Later*).

This case is not, however, a counterexample to the TRIA. Harman believes it is because she thinks that the TRIA implies that it is better not to perform the surgery. This is because it does not discount the immediate, lesser suffering that the surgery would cause but greatly discounts the greater suffering that will result if the surgery is not performed, since the horse is now only weakly

⁸ I discussed this problem in section 2.4 of chapter 5 in *The Ethics of Killing*, but without achieving any resolution.

⁹ Elizabeth Harman, "The Moral Significance of Animal Suffering and Animal Death," in Tom L. Beauchamp and R. L. Frey, eds., *The Oxford Handbook on Ethics and Animals*, (New York: Oxford University Press, 2011), p. 735.

psychologically connected to itself at the later time when the greater suffering would occur. But hers is a case in which the horse will exist in both possible outcomes of one's choice. Over time, its interests include a strong present interest in avoiding the immediate suffering, a weak present interest in avoiding the later suffering, and a strong future interest in avoiding the later suffering. Although the TRIA discounts interests for psychological unconnectedness, it does not claim that only present interests matter; rather, it accepts that all the interests that an individual has at all times at which it exists can be sources of reasons. Thus, an agent who has no reason to care more about the horse now than in the future should choose to treat it because its *future* interest in avoiding the greater suffering will be stronger than its present interest in avoiding the immediate, lesser suffering. (There is a parallel here with a relevant difference between abortion and prenatal injury. Whereas abortion frustrates only the weak present interest of a fetus in continuing to live, prenatal injury may frustrate the strong future interests of a person.)

In *Suffering Now* and *Suffering Later*, however, one of the options is that the animal will die now. Because of this, the only interests that it has independently of the outcome of one's choice are its present interests. In *Suffering Now*, the horse now has a strong interest in avoiding immediate suffering but only a weaker interest (because of the weakness of the psychological connections between itself now and itself later) in experiencing greater happiness in the distant future. Its interests over time do *not* include a strong interest in experiencing greater happiness in the distant future that will exist at that later time independently of the decision about whether to save it. Similarly, in *Suffering Later*, it has a strong interest in having happiness now but only a weaker interest in avoiding greater suffering in the distant future. Whether it will later have a strong interest in avoiding great suffering at that future time depends on the decision about whether to save it. In short, in Harman's example, the strong interest that the horse will have in avoiding great suffering in the future is an actual, future interest that one must take into account in determining what is best for the horse. In *Suffering Now*, by contrast, the strong interest the animal might have in experiencing greater happiness in the future is only a *possible* interest, as is the strong interest it might have in avoiding greater suffering in the future in *Suffering Later*. Thus, while Harman's case is not a counterexample to the TRIA, the related case of *Suffering Later* is; and *Suffering Now* may be as well.¹⁰

5 Possible Interests

¹⁰ Harman might argue that, although her example is presented as a choice between doing and not doing the surgery, so that the horse's later interest in avoiding suffering is a *future* interest relative to that choice, there is nothing that excludes killing the horse as an option. If we explicitly include that among the options, the horse's later interest becomes a *possible* interest and her example then has the right form for a counterexample to the TRIA.

Those who believe Suffering is Worse should accept the TRIA's judgment that it is better to allow the animal to die in Suffering Now. But they and others will reject its implication that it is better to save the animal in Suffering Later, when its future life would contain significantly more suffering than happiness. It seems that the relevant interests in Suffering Later include not only the animal's undiscounted present interest in experiencing immediate happiness and its discounted present interest in avoiding greater suffering in the future, but also the possible interest it might have later in avoiding the greater suffering when that suffering might occur. This later interest would be strong, as it would be an interest in the animal's immediate experience and thus would not be discounted for psychological unconnectedness. If this *possible* interest can ground a reason now to prevent the animal from experiencing the greater suffering, that could explain and justify our belief that it would be better not to save the animal's life.

There are, however, certain problems with this. If an interest an animal would have only if we act in a certain way can give us a reason either to act or not to act in that way, then the possible interest the animal in Suffering Now might later have in experiencing happiness in the distant future could ground a strong reason now to enable it to experience that happiness. That possible interest could outweigh the animal's undiscounted present interest in avoiding the immediate suffering. But if the TRIA were to take account of possible interests in this way, it would then imply that it would be best to save the animal in Suffering Now. It would no longer support the view that Suffering is Worse.

Many will find this implication plausible. But there is more at stake than whether Suffering is Worse is true. As I noted earlier, the TRIA offers a plausible explanation of why death is worse for a 20-year-old than for a fetus, even though the fetus loses more good life in dying. If, moreover, a fetus lacks the moral status that would give it a right not to be killed that is independent of the strength of its interest in continuing to live, the TRIA's explanation of why a fetus is not greatly harmed by being killed provides the basis for an argument for the permissibility of abortion in a wide range of cases.¹¹ Abortion can be permissible when the weak interest of the fetus in continuing to live can be outweighed by conflicting interests of the pregnant woman. But if possible interests can ground reasons now (as an animal's possible interest in not experiencing great suffering in the future seems to ground a reason not to save its life now), then a fetus's possible interests in having the good experiences of its later life can ground a reason not to kill it that is independent of any actual interest it has now. In the remainder of this article, I will consider whether the TRIA can be supplemented in a way that enables it to address the two counterexamples (Suffering Now and Suffering Later) without forfeiting its ability to explain why a very early death is a lesser misfortune and why abortion can often be permissible.

6 Unconnected Animals and the Asymmetry

¹¹ This argument is developed in chapter 4 of *The Ethics of Killing*

In section 3, I suggested that the death of an individual immediately after it begins to exist is of scarcely more significance than its not having come into existence at all would have been, given the absence of psychological connections between the individual at the time of death and itself as it would have been when the good things in its life would have occurred. I claimed this of persons but it should be true of animals as well. Consider now an animal that lacks any degree of self-consciousness and thus, if it lives, will be psychologically unconnected to itself beyond the immediate future. Call such an animal a *psychologically unconnected animal*, or “unconnected animal” for brevity. Because the relation between itself now and itself in the future is relevantly like that of an animal that has just begun to exist, it seems that it would suffer no greater misfortune in dying than an animal with an equal amount of good life in prospect that has just begun to exist. It thus seems that whether this unconnected animal continues to exist matters no more, even for its own sake, than whether an animal that would have a similar future comes into existence.¹²

The argument about abortion to which I referred depends on a similar claim. Even a fetus with the capacity for consciousness is at most only very weakly psychologically related to itself as it might later be as a person. Whether it continues to exist is therefore relevantly more like whether a person comes into existence than it is like whether an existing person, who would be strongly psychologically connected to herself in the future, continues to exist.

If the continuing to exist of an unconnected animal is relevantly like the coming into existence of a similar animal, we may be able to draw inferences about the morality of saving or killing such animals from our views about causing them to exist or preventing them from existing. And what we should believe about causing animals to exist should parallel what we should believe about causing people to exist. The view that most people seem to accept is what is sometimes called the *Asymmetry*, the view that while the expectation that a person would have a life that would be bad for him grounds a moral reason not to cause or allow him to exist, the expectation that a person would have a life that would be good for him does not ground a reason to cause or allow him to exist. Suppose we understand the *Asymmetry* to include the view that the reason not to cause a person to exist whose life would be bad for him is as strong as the reason not to cause an existing person’s life to be bad to a roughly equivalent degree. So understood, the *Asymmetry* can also apply to causing animals to exist.

(One might wonder whether it is consistent to accept this *Asymmetry* but to reject, as I did earlier, a quite general asymmetry between suffering and happiness. My objection to the general asymmetry was that it seems incompatible with the permissibility of causing a person great suffering to save her life, even when it is not possible to get her consent, provided the net benefits

¹² This claim is similar to that which is the basis for Peter Singer’s suggestion that non-self-conscious animals may be “replaceable.”

of further life would outweigh the suffering. There would thus be an inconsistency if the Asymmetry implied that it would be impermissible to cause an unconnected animal great suffering to enable it to continue to live, when the benefits of continued life would outweigh the suffering. And I have suggested that the Asymmetry does have implications for causing or allowing unconnected animals to continue to live, on the assumption that a wholly unconnected animal's continuing to live is relevantly like a similar animal's coming into existence. But even on that assumption, the Asymmetry does not object to causing an unconnected animal to suffer to enable it to continue to live, provided the suffering would be outweighed by the net benefits of continued life. What the Asymmetry denies is that the benefits an individual would get from life provide a reason to cause that individual to exist or, in the case of an unconnected animal, a reason to cause it to continue to exist. It does not deny that these benefits can weigh against and offset the suffering the individual might also experience. One might express this by saying that the Asymmetry denies that pure benefits have "reason-giving weight" but not that they have "offsetting weight."¹³)

We now have the elements of an argument that explains how an unconnected animal's possible future interests are relevant to whether it is better for it to be caused or allowed to continue to exist or instead caused or allowed to die.

- (1) Whether an unconnected animal continues to exist is relevantly like whether an animal with similar prospects comes into existence.
- (2) The Asymmetry: if an animal's life would be intrinsically bad for it, there is a moral reason not to cause or allow it to come into existence. But if its life would be intrinsically good for it, that does not ground a reason to cause or allow it to come into existence.
- (3) If, therefore, an existing unconnected animal's future life would be bad for it, there is a moral reason to kill it or to allow it to die. But if its future life would be good for it, that alone does not constitute a reason to cause or allow it to continue to exist.

According to this argument, there is a moral asymmetry between causing or allowing an unconnected animal to continue to exist when its future life would be bad and causing or allowing it to continue to exist when its life would be good – just as there is thought to be a parallel asymmetry in causing or allowing individuals, whether persons or animals, to exist. In both types of case, there is,

¹³ For elucidation of this distinction, see Jeff McMahan, "Causing People to Exist and Saving People's Lives," *Journal of Ethics* 17 (2013): 5-35. There I use the term "canceling weight" rather than "offsetting weight," though I now think the latter is preferable.

at the time of choice, no individual that has a present (or future) interest in the good or bad things that might occur in what might be its future life. There are only the *possible* interests that either a possible animal or an unconnected animal might later have in avoiding suffering or experiencing happiness. (Recall that the reason why Suffering Later is an effective counterexample to the TRIA only its own is that the animal has no present or future interest in its own possible suffering in the distant future.) But assuming the Asymmetry applies in both types of case, possible interests in avoiding suffering count the way present interests in avoiding suffering count, though possible interests in experiencing happiness do not count at all – *except* in weighing against and offsetting possible interests in avoiding suffering.

That final qualification is important. We accept that the reason not to cause or allow a possible interest in avoiding suffering to exist and be frustrated can be overridden if the individual who would have that frustrated interest would be more than compensated by the existence and satisfaction of possible interests in later experiencing greater happiness. We accept, in other words, that while possible happiness that no individual has a present interest in having may provide little or no reason to cause an individual to exist or to continue to exist (that is, it may have no reason-giving weight), it can weigh against and potentially offset a comparable amount of suffering the individual might later experience (that is, it has offsetting weight). For if this happiness had no offsetting weight, there would be a presumption against causing or allowing individuals to exist, as well as against causing or allowing unconnected animals (or unconnected fetuses) to continue to exist, even when their future lives would be overall worth living. I take this to be a *reductio* of the suggestion that happiness has no offsetting weight in these types of choice.

One might think that if the TRIA is conjoined with the Asymmetry, the permissibility of causing or allowing an unconnected animal to exist or to continue to live depends on the order in which the good and bad elements of its future life would occur, as it does in Suffering Now and Suffering Later. But in fact the order is irrelevant. In the absence of psychological connections between the animal now and itself in the future, the animal is equally unrelated to all parts of its future and there is no basis for discounting some relative to others.

7 Connected Animals and Weak Asymmetries

While some animals do seem to live entirely in the present, others are psychologically connected to themselves in the future in relevant ways. We can call such animals *psychologically connected animals*, or “connected animals” for brevity. Such animals are, of course, psychologically connected to themselves in the future to varying degrees. The nonhuman great apes, for example, are more closely connected than other animals that have a lower degree of self-consciousness. But even the more highly connected animals are now only weakly psychologically connected to themselves in the future in comparison with the degree to which persons are connected to their future selves. Still, if the Asymmetry is correct, the comparative weakness of a connected animal’s present interest in avoiding future suffering does not weaken the reason to prevent that

suffering. In particular, an animal's being prevented from having a future life that would be on balance bad for it matters equally whether the animal is unconnected or connected – that is, whether or to what extent it now would be psychologically related to itself in the future. Indeed, preventing an animal from suffering in the future matters equally even if the animal does not at present exist and may never exist.

Matters are different, though, when we consider the prospect of future happiness. Because a connected animal would be psychologically related to itself in the future, it has a present interest in experiencing future happiness and there is thus a dimension to the badness of its loss of good life in the future that is not present in the equivalent loss that occurs when an animal fails to come into existence or when an unconnected animal dies. Its loss of good life through death thus matters more than an equivalent loss by an unconnected animal, and the stronger the psychological relations between the animal now and itself in the future would be, the more its loss matters – that is, the greater its misfortune would be in suffering that loss. Thus, in the case of a connected animal, though not in the case of an unconnected animal, the order in which the good and bad elements of its future life would occur may matter. Assuming, for example, that the strength of a connected animal's psychological connections with itself in the future would diminish with time, the loss of later good life matters less than the loss of earlier life that would be equally good.

In summary, if the Asymmetry is correct, the expectation that an animal, whether connected or unconnected, would have a life worth living provides no reason to cause it to exist. Nor does that expectation ground a reason to cause or allow an unconnected animal to *continue* to live. But the expectation that a connected animal would have a life worth living does provide some reason to cause or allow it to continue to live.

Some philosophers have, however, challenged the Asymmetry and argued instead for a *Weak Asymmetry*, according to which the expectation that an individual's life would be on balance intrinsically good – or good beyond some minimum level of goodness – does ground a reason to cause or allow that individual to exist, though one that is weaker than the reason not to cause or allow an individual to exist whose life would be intrinsically bad to a roughly equivalent extent.¹⁴ Assuming that a unconnected animal's continuing to live is relevantly like the coming into existence of an animal with similar prospects, a Weak Asymmetry also implies that an unconnected animal's possible later interest in experiencing happiness provides some reason to cause or allow it to *continue* to live, though not as strong a reason as its possible later interest in avoiding a comparable amount of suffering would provide for *not* causing or allowing it to continue to live.

¹⁴ See, for example, *ibid.*, and Elizabeth Harman, "Can We Harm and Benefit in Creating?" *Philosophical Perspectives* 18 (2004): 89-113.

(One question about which I am uncertain is whether, if we accept a Weak Asymmetry, we should also accept that the reason to cause or allow a *connected* animal to continue to live is stronger than it would be if the Asymmetry were true. It may seem that the reason implied by a Weak Asymmetry to cause or allow an unconnected animal to continue to live applies as well to a connected animal and combines additively with the reason to extend its life that derives from its psychological connectedness to its future selves. But, if this is true, it seems that the reason to extend the life of a *person* should also be stronger if a Weak Asymmetry is true than it would be if the Asymmetry were true. But it does not seem that the strength of the reason not to kill people, or to save them, depends upon whether there is a moral reason to cause people to exist if their lives would be worth living. I leave this issue open here.)

8 Implications of the TRIA Combined with the Asymmetry

Suppose we combine either the Asymmetry or a Weak Asymmetry with the TRIA. The resulting views have mostly plausible implications for the cases that challenged the TRIA on its own – namely, Suffering Now and Suffering Later. Let us consider instances of Suffering Now and Suffering Later involving different kinds of animal and apply the different principles to them. Consider first the conjunction of the TRIA and the Asymmetry.

(1) *Suffering Now with an unconnected animal.* The animal has little present interest in avoiding the imminent suffering and no present interest in experiencing greater happiness in the distant future. According to the Asymmetry, the suffering it would experience grounds a reason to prevent it from continuing to live while the greater happiness it might experience provides no reason to enable it to continue to live. It may therefore seem that it would be better not to save it. While this is consistent with the view that Suffering is Worse, it is intuitively implausible. But while the happiness the animal might experience has no reason-giving weight, it nevertheless has offsetting weight vis-à-vis the suffering that would occur earlier. Because there are no psychological relations between the animal now and itself at any time beyond the immediate future, the offsetting weight of the possible future happiness cannot be discounted for psychological unconnectedness. The temporal order of the suffering and happiness does not affect their respective weights. Hence the prospect of the greater happiness offsets the lesser suffering, making it permissible to save the animal. But because the happiness has no reason-giving weight, it is also permissible to allow the animal to die, just as it would be permissible not to cause a similar animal to exist. This seems a plausible result.

(2) *Suffering Later with an unconnected animal.* The animal has little present interest in experiencing the immediate happiness and no present interest in avoiding the greater suffering later. But according to the Asymmetry, there is a reason to prevent an individual from suffering even if it has no present interest in avoiding it. This reason is as strong as the reason to prevent an existing individual of the same sort from experiencing equivalent suffering now. It is therefore better, on this combination of views, not to save the animal, which is intuitively the correct result.

(3) *Suffering Now with a connected animal.* This seems to me the most problematic case. In cases involving causing individuals to exist, or causing unconnected animals to continue to live, there can be no discounting of the offsetting weight of future happiness for psychological unconnectedness. But a connected animal would be more closely psychologically connected to itself in the immediate future than in the distant future. It therefore has at least a moderately strong present interest in avoiding the immediate suffering but only a very weak present interest in having the greater happiness in the distant future. According to the Asymmetry, there is no reason independent of the animal's present interests to enable it to have the greater happiness for its own sake. The question is whether the greater happiness in the distant future can offset the lesser suffering in the immediate future. It seems, however, that the offsetting weight that the animal's possible later happiness has vis-à-vis its possible immediate suffering ought to be discounted for the weakness of the psychological connections between itself now and itself later. Even though the happiness it might experience later is substantially greater than the suffering it might endure now, it may be that the animal cannot be compensated for lesser suffering in the immediate future by the greater happiness of a later individual that would be so weakly psychologically related to itself now, even though that individual would be itself. It may be, in other words, that the animal's steeply discounted interest in experiencing the greater happiness cannot outweigh its only slightly discounted interest in avoiding the lesser suffering. If so, it may be better to allow the animal to die, even though its future life would be on the whole worth living. This, I concede, seems implausible, but it is not highly implausible, and it does provide limited support for Suffering is Worse, which many people accept.

(4) *Suffering Later with a connected animal.* The animal has a moderately strong present interest in experiencing the immediate happiness but only a very weak present interest in avoiding the distant greater suffering. But according to the Asymmetry, there is a reason to prevent the suffering that is independent of any present interest in its avoidance. Because the suffering outweighs the happiness, it is better to allow the animal to die.

9 Implications of the TRIA Combined with a Weak Asymmetry

Next consider the implications for the same four cases of a combination of the TRIA and a Weak Asymmetry.

(5) *Suffering Now with an unconnected animal.* In the case of an unconnected animal, the order in which good and bad experiences occur makes no difference because there are no psychological relations that could be either stronger or weaker over time. According to any Weak Asymmetry as understood here, the suffering in the immediate future counts fully, independently of any actual interest in its avoidance, while the prospect of greater happiness in the distant future provides only a weaker reason to enable the animal to continue to live. But even if a Weak Asymmetry gives far greater weight to the avoidance of suffering than to the experience of happiness, these are reason-giving weights, so that all the Weak Asymmetry implies is that there may be no duty, all things

considered, to save the animal *for its own sake now*. It allows that the greater happiness has sufficient offsetting weight to make it permissible to save the animal.

(6) *Suffering Later with an unconnected animal*. A Weak Asymmetry implies that the happiness in the immediate future provides some reason to enable the animal to continue to live. But both because the greater suffering is not discounted for psychological unconnectedness over time and because that suffering has full offsetting weight vis-à-vis the happiness, it is better to allow the animal to die.

(7) *Suffering Now with a connected animal*. The remarks under 3 above apply here as well. Whether it would make a difference if a Weak Asymmetry were true depends on the question raised parenthetically at the end of section 7 – namely, whether the reason implied by a Weak Asymmetry to cause an individual to exist if its life would be worth living combines additively with a connected animal's psychological connectedness to itself in the future to strengthen the overall reason to cause or allow it to continue to live. If it does, it is more likely that the combined TRIA and Weak Asymmetry will imply that it is better to cause the animal to survive by treating it. And that would be a reason to accept a Weak Asymmetry rather than the Asymmetry. But, as I noted, it is uncertain whether a reason to cause individuals to exist would then be an independent reason to keep them in existence in addition to the familiar reasons concerned with interests and rights to cause or allow connected animals and persons to continue to live.

(8) *Suffering Later with a connected animal*. Assuming a Weak Asymmetry, the animal's present interest in experiencing the immediate happiness provides a reason to enable it to continue to live that is stronger than the reason to enable an unconnected animal to continue to live. But we are assuming that a Weak Asymmetry also implies that the reason to prevent the possible later suffering is as strong as the reason there would be to prevent the animal from experiencing equal suffering now. So it is better to allow the animal to die.

These implications of the two views for the four possible cases are mainly plausible, with 3 and 7 as exceptions – though those who find Suffering is Worse plausible may welcome the implications described in 3 and 7. What seemed to be implausible implications of the TRIA for unconnected animals in Suffering Now are blocked by appeal to the distinction between reason-giving weight and offsetting weight, while the TRIA's implausible implications for Suffering Later are blocked by combining it with either the Asymmetry or a Weak Asymmetry. There are, of course, objections to these latter views, but abandoning them in favor of full Symmetry between causing miserable individuals to exist and causing well-off individuals to exist would have highly counterintuitive implications.¹⁵

¹⁵ McMahan, "Causing People to Exist."

10 Abortion

In this short section I will consider the implications of the two views (the TRIA combined with the Asymmetry and the TRIA combined with a Weak Asymmetry) for the morality of abortion. One could plausibly argue that a human fetus is relevantly like an unconnected animal, in that it would be wholly psychologically unrelated to itself in the future. But I will assume that a fetus is instead relevantly like a connected animal and would be weakly psychologically related to itself as a person. This assumption is less favorable to the view that abortion is often permissible. If a fetus is relevantly like a connected animal, then both views imply that if a fetus faces the prospect described in *Suffering Now*, it could be better to allow it to die. That, I have conceded, is implausible, though perhaps not highly implausible.

It is, however, improbable that a fetus could face a prospect such as that in *Suffering Now*. The typical expectation is that a fetus's future life would be overall worth living, with good and bad elements more or less evenly distributed throughout. That is, good and bad experiences normally alternate, with more good experiences than bad. Suppose we accept the TRIA and the Asymmetry. Even though the suffering the fetus would experience at various times in its life counts against allowing it to continue to live, the greater happiness it would experience *around the same later times* outweighs that suffering and thus offsets the reason to prevent it from continuing to live. This is so despite the fact that the fetus now would be only weakly related to itself during the periods of happiness. Suppose that a fetus would later, as a person, experience suffering at a certain time, but would also experience happiness shortly before and shortly after the suffering. Given that the fetus now would be psychologically related to itself to much the same degree at all three times, both the happiness and the suffering should be discounted to the same degree in the determination of their offsetting weights. Thus, if the experiences of happiness would together be greater in amount than the suffering, they would offset it. If, therefore, the fetus's life would be worth living overall, it would be permissible to allow it to continue to live. But given the Asymmetry and the fact that the fetus would be only very weakly psychologically related to itself as a person, the happiness it might later experience has very little reason-giving weight – that is, it provides only a very weak reason to cause or allow the fetus to continue to live. The fetus's weak interest in continuing to live could therefore be outweighed by the interests of the pregnant woman.

It is, of course, insufficient to deal with a counterexample to show that the case is unlikely to arise in practice. But there is more that can be said. The TRIA asserts that the misfortune of death is a function of two factors: the value of the life lost and the degree to which the victim would have been psychologically connected to that life. It does not assert what the relative weights of the two factors are. It is therefore possible that, even though a fetus would, in a case like *Suffering Now*, be much more closely psychologically connected to itself in the near future than in the distant future, the sheer magnitude of the value of the later life it would lose in dying could outweigh the immediate suffering. This is far more likely to be true in the case of a fetus than in the case of a connected

animal because persons generally have lives of much greater quality and duration than those of animals. It is also possible that, even though the fetus's present interest in having greater happiness later outweighs its interest in avoiding suffering now, its overall interest in continuing to live is still sufficiently weak to be outweighed by the interest of the pregnant woman in having an abortion.

11 Conclusion

In the last few sections I have sought to defend the TRIA against objections by combining it with either the Asymmetry or a Weak Asymmetry. If it can be defended in this way, it may provide some support for Humane Omnivorism, as it implies that death is a lesser misfortune for animals than for persons not only because their future lives would be less good but also because of their lesser psychological connectedness to themselves in the future. In particular, if the TRIA is correct, the painless killing of wholly unconnected animals is relevantly like preventing animals with comparable prospects from coming into existence, which few believe would be wrong.

The reason not to kill connected animals is different and stronger, but less strong than the reason not to kill persons, at least in part because their loss of a less valuable future to which they would be less closely connected is a lesser misfortune. Still, even if the TRIA is true, whether Humane Omnivorism is permissible depends on several considerations. One is whether animals used for food are unconnected or connected. This is of course an empirical matter but I suspect that only the really lower forms of animal are wholly unconnected. Most of the animals that could be humanely reared for human consumption are connected to varying degrees: pigs more than cows and cows more than chickens. And the relevant connections are stronger between the animal now and itself in the near future than between itself now and itself in the further future. This makes it possible that the satisfaction that a connected animal would get just from eating over, say, the next month could outweigh the difference in pleasure that people would derive from eating its meat rather than eating vegetarian meals. This too is an empirical question but again it seems doubtful that in most cases the difference in human pleasure would outweigh the loss of animal pleasure. If, moreover, we reject the TRIA, so that none of the happiness that an animal might experience over many further years of life would be discounted, the great majority of killings required by Humane Omnivorism would inflict on the animal victim a loss that would be vastly greater than the benefits it would provide to "humane" omnivores.¹⁶

¹⁶ For written comments on an earlier draft, I am deeply grateful to Roger Crisp, Maxime Lupoutre, Theron Pummer, Tatjana Visak, and especially Derek Parfit.