

*Eating animals the nice way*

Many people are opposed to factory farming because of the terrible suffering it inflicts on animals, yet see no objection to eating animals that are killed painlessly after having been reared in conditions that are at least no worse, and are perhaps even better, than typical conditions in the wild. Let us refer to this latter practice, in which animals are reared for human consumption but in humane conditions, as ‘benign carnivorousness.’ When philosophers discuss the morality of this practice, they sometimes argue that, unlike animals killed by hunters, animals that are raised to be killed and eaten would never have existed if we had not created them in order to eat them. If benign carnivorousness enables these animals to have contented lives that they would otherwise not have had, it seems better for the animals as well as for the people who get to eat them. How, then, could such a practice be objectionable?

Those who object to eating factory-farmed animals but accept benign car-

---

*Jeff McMahan is Professor of Philosophy at Rutgers University and author of “The Ethics of Killing: Problems at the Margins of Life” (2002).*

---

© 2008 by the American Academy of Arts & Sciences

nivorism generally believe that while animal suffering matters, animal lives do not – or at least not as much. They think that there is a strong moral reason not to cause animals to suffer, and even to try to prevent them from suffering, but not a comparably strong reason not to kill them, or to ensure that they have longer rather than shorter lives.

One possible basis for this view is the difference between how well off and how badly off it is possible for animals to be. Although animals are incapable of the depths of psychological misery to which most human beings are susceptible, their capacity for physical suffering rivals our own. Yet their highest peaks of well-being are significantly lower than those accessible to most human beings. While some animals – dogs, for instance – experience exuberant joy more readily and frequently than many adult human beings do, animals lack other dimensions of well-being that are arguably more important, such as achievement, creativity, deep personal relations, knowledge, aesthetic appreciation, and so on.

There is another, possibly even more important, reason why animal lives matter less than animal suffering. Not only do animals’ future lives promise less in terms of both quality and quantity of

good than those of most human beings, but animals are also less strongly connected to themselves in the future in the ways that make it rational to be concerned about an individual's future well-being for that individual's own sake now. Because they are not self-conscious, or are self-conscious only to a rudimentary degree, they are incapable of contemplating or caring about anything more than the immediate future. They do not, therefore, have desires or intentions or ambitions for the future that would be frustrated by death.<sup>1</sup>

Yet the lives of animals must matter to some extent – that is, animals must have an interest in living to experience the goods that lie in prospect for them. In particular, the goods that an animal's future life could contain must matter enough to justify allowing the animal to endure a certain, even considerable, amount of suffering. For if an animal's avoidance of suffering were significantly more important than its living to experience the goods that its future life could contain, then it would be better *for the animal* to be painlessly killed before it could undergo any suffering at all.

But this is implausible. It can be better for an animal to endure a certain amount of suffering if the good experiences it might have afterward would be sufficient to outweigh the suffering. We all acknowledge this when we submit our pets – just as we submit ourselves – to painful but life-saving medical treatments.

The upshot of these reflections is that there is reason to be skeptical of the widespread view that the prevention of

suffering among animals is much more important than the extension of their lives. This is not to deny that there is a significant difference between persons and animals in this respect. The goods that are characteristic of human life are so much higher than those characteristic of animal life that it is rational for us to tolerate substantially more suffering in order to continue to live than it would be acceptable to make an animal endure in order to save its life. But the goods of an animal's life weigh against the evils in the same way that goods and evils weigh against one another in the life of a person. It is just that animal goods are lesser goods, and therefore have less weight.

According to some advocates of benign carnivorousism, it is precisely because the lives of animals raised in humane conditions are good that the practice is not only permissible but desirable. If the lives the animals have are good, and if they would not have existed at all without the practice, then at the very least benign carnivorousism cannot be worse for them. And since eating animals that have been humanely raised and painlessly killed may be better for people than having to go without meat altogether, the practice would be, at a minimum, better for some and worse for none – or, as economists say, Pareto optimal. But it may even be better for everyone affected, animals included. (Here I ignore the larger question of whether meat-eating is worse for people because it involves an inefficient use of the world's resources.)

While the case for benign carnivorousism is often stated this way, these claims are misleading. The claim that benign carnivorousism would not be worse for the animals that it would cause to exist is, strictly speaking, trivially true, while the claim that it would be better for them is necessarily false. This is because 'worse'

1 For discussion of the relevance of psychological continuity within a life to the ethics of killing, see Jeff McMahan, *The Ethics of Killing: Problems at the Margins of Life* (New York: Oxford University Press, 2002), 39 – 43, 69 – 82.

and 'better' are comparative terms, and one element in each implied comparison is never existing at all.

Consider the claim that it is not worse for an animal to be caused to exist. This is not a substantive claim. It is instead true as a matter of logic, since it is *incoherent* to suppose that an animal's being caused to exist could be *worse* for it. Because 'worse for' is comparative, the claim that it is worse for an individual to be caused to exist implies that it would have been *better for that individual* not to have been caused to exist – that is, never to have existed at all. But there cannot be anyone for whom it is better never to exist.

Similarly, to say that it is better for an animal to be caused to exist implies that it would have been worse for that same animal never to have existed. But again, there cannot be anyone for whom it is worse never to exist. In one clear and relevant sense, there *are* no individuals who never exist.

It is thus true, even of an animal whose life involves nothing but unrelieved agony, that it is not *worse* for it to exist.<sup>2</sup> It can certainly be *bad* for that animal to exist, and to have been caused to exist. 'Bad' is noncomparative. We can say that a life is bad if its bad aspects outweigh the good. And it can be bad for an animal to be caused to exist with a life that is bad – as is generally the case of animals that are factory farmed.

Just as it can be bad to be caused to exist with a life that is bad, so it can be good to be caused to exist with a life that is good. Since benign carnivorousness by

2 Or, rather, not worse for it than never to exist. It does seem that to exist can be better or worse for an individual than to cease to exist. Contrary to what Epicurus once claimed, we can make sense of the idea that there is someone for whom ceasing to exist is worse, or better, than continuing to exist.

definition aims to cause animals to exist with lives that are good – in which the good elements outweigh the bad – it is plausible to say that the practice is good for the animals it causes to exist, even if the ultimate aim is to make them available for human consumption. While the practice also involves painlessly killing them, and while killing them is bad for them, and worse for them than allowing them to continue to live, the practice as a whole is still good for them, since their lives are good and otherwise they would not have existed at all.

Benign carnivorousness is, moreover, a continuing practice. When some animals are painlessly killed for consumption, others are caused to exist in their place. The practice thus yields a continuous bounty of contented animals and contented diners.

Before considering some objections to benign carnivorousness, we should pause to summarize and review the ideal conditions of the practice.

- The animals would have lives worth living. They would be well fed, protected from predators, allowed the free exercise of their natural instincts, and at least as well-off overall as their counterparts living in the wild.
- They would not have existed if not for the practice of benign carnivorousness. Moreover, it is not just that the *particular* animals would not otherwise have existed; it is that far *fewer* animals with lives worth living would have existed in the absence of the practice.
- The animals would be allowed to live a considerable portion of their natural life span before being painlessly killed.
- Although killing the animals might deprive them of several years of life, the amount of good they would thereby lose is comparatively slight.

- The significance of the loss the animals suffer must be discounted for the relative absence of psychological unity in their lives.
- Those that are painlessly killed are replaced by new animals with lives that are equally good.
- The pleasure that people get from eating the animals is in general greater than the pleasure these people would have gotten from eating foods derived entirely from plants.

The question now is whether a practice that has these features, or at least many of them, is morally permissible.

One obvious point is that no one would invoke the logic of the argument just given to justify a parallel practice involving *persons*. Imagine that the world's population has reached a point at which people have agreed to adopt a policy of replacement – that is, people may have a child only when someone dies, so that total population does not increase. Suppose further, however, that there remains a chronic shortage of donor organs and that many people continue to die for want of an organ transplant. In these conditions, people might agree to allow a certain number of people to be born above the limit, provided that they will be painlessly killed at the age of fifty in order to make their organs available for transplantation. Even though these people would have lives well worth living and would never have existed had we not caused them to exist to be able to use their organs, and even if the benefits to the recipients of their organs would be significantly greater than the harm the victims would suffer (perhaps because their organs would be given only to recipients under the age of thirty), this practice would clearly be wrong.

It would be wrong presumably because persons have rights that constrain others from using them in certain harmful ways even when the practice that involves using them in these ways would not be bad for them, and might even be good for them overall. It would not matter that we had brought these people into existence only on the condition that we could kill them at the age of fifty. Once they become persons, they have a right not to be killed. It would be irrelevant that it was good for them to exist and that they would never have existed had we not caused them to exist specifically in order to kill them for their organs.

If animals had the same rights as persons, those rights would provide a decisive objection to benign carnivorousness. But it is hard to believe that killing an animal is morally objectionable for the same reasons and to the same degree as killing a person. Of course, human intuitions about the moral status of animals are so contaminated by self-interest and irrational religious belief as to be almost wholly unreliable. Yet even most people who have become vegetarians or vegans for moral reasons would accept the permissibility of killing an animal if what was at stake were as important as saving the life of a person. This would be true even if the animal were one of the higher primates. Suppose, for example, that the painless killing of a single chimpanzee could save the lives of two five-year-old children by making its organs available for transplantation. Although virtually no one believes that it could be permissible to kill one five-year-old child in order to use her organs to save two other five-year-olds, most of us believe that it would be permissible to kill the chimpanzee, and could produce arguments to show that this belief is not speciesist but is based on morally significant intrinsic differences between chimpanzees and

*Eating animals the nice way*

normal five-year-old human beings. But if this robust intuition that xenotransplantation can be permissible is right, then animals do not have the rights that we, as persons, have.

It might be that animals have rights, but rights that are weaker than those of persons. If so, an animal's rights might be overridden when it is necessary to kill it to save the life of a person, but not when the only purpose that would be served by killing it is to enable someone to enjoy the taste of meat. This view is, however, hard to reconcile with the nature of rights. A right of a given type – in this case, a right not to be killed – is generally held to be invariant in strength among all those who possess it. A theory that allowed the strength of rights of a given type to vary with the strength of the interests they protect would hardly differ in substance from a theory enjoining respect for interests.

It might be, of course, that we attribute equal rights to all human beings in order to articulate a conception of human equality. And it is compatible with human equality that animals could have weaker rights of variable strength. But this view could be true only if species membership were relevant to the possession of rights, which I have argued elsewhere is not the case.<sup>3</sup>

So if, as I believe, xenotransplantation could be permissible, it seems that the explanation of why it may be wrong to harm or kill animals for lesser reasons derives from a requirement of respect for their *interests*. An appeal to rights is necessary only when a principle requiring respect for interests cannot account for the moral reasons we seem intuitively to have. And these reasons seem to arise only in our dealings with individuals that have not only interests but also

certain higher cognitive and emotional capacities, such as self-consciousness, autonomy, and rationality.

If I am right that animals do not have the rights that protect persons from certain forms of harmful using, we cannot reject benign carnivorousness on the grounds on which we would rightly reject a practice that would cause *people* to exist in order to use them later in harmful ways.

Suppose, then, we consider benign carnivorousness in terms of the interests at stake. Consider an animal whose flesh could provide one meal each for twenty people. How might the human and animal interests compare? It seems that we have to compare the animal's interest in continuing to live – a function of both the amount of good that its life would contain were it not killed, and the degree to which it would be psychologically connected to itself in the future – with twenty people's interests in the pleasure they would get from eating the animal.

It is important to stress that the people's interest is *not* in having the pleasure of eating meat rather than having no pleasure at all; it is instead in the *difference* in pleasure between eating meat and eating food derived from plants. Given comparable investments in the procurement and preparation of the two types of food, this difference is likely to be slight. Note also that the time that a person spends tasting meat during a normal meal is not much longer than a few minutes. It therefore seems unlikely that the interests that twenty people each have in experiencing a few minutes of slightly greater pleasure could outweigh all the good that an animal's life might contain over several years, even when that good is heavily discounted for the absence of significant psychological continuity within the animal's life.

<sup>3</sup> See McMahan, *The Ethics of Killing*, 203–217.

Some people will no doubt think: 'How typical, and predictable, that an academic philosopher would scorn, or affect to scorn, the pleasures of eating. For most people, the pleasures of eating, particularly in a social context, are among the great goods of human life.'

But those who press this point undermine their own case. It does seem that, for many people, meals and snacks are among the few intervals of pleasure that enliven their otherwise quotidian lives. Yet anyone who has ever lived with dogs, horses, or other animals knows that many animals also take great pleasure in eating. There is a reason why eating is often referred to as an 'animal pleasure,' in contrast, for example, to the pleasure of listening to a symphony. Thus, if we add up the differences in pleasure that twenty people would get at one meal from eating meat rather than food derived from plants, and compare that total pleasure with the pleasures that the animal would get from several years of eating several times a day (not to mention the other pleasures its life would contain), it is scarcely credible to suppose that the people's interests could outweigh those of the animal.

It may seem that we have lost sight of the important point I highlighted earlier: that the animals that would be eaten in a practice of benign carnivorousness would owe their existence to the practice. They would have many meals, and therefore much pleasure, but only if people were to bring them into existence in order to eat them. Surely, one might argue, we ought to take this fact into consideration in assessing how the practice of benign carnivorousness bears on both human and animal interests.

There are, however, no animal interests that favor instituting a practice of benign carnivorousness. No individual, ani-

mal or otherwise, has an interest in being caused to exist. Interests arise only once an individual exists; therefore, to cause an individual to exist cannot be to satisfy any interest of that individual. It may be *good* for animals to be caused to exist by the practice of benign carnivorousness; but that is compatible with there being no reason to have the practice that is grounded in animals' *interests*.

If, therefore, we evaluate the practice of benign carnivorousness by reference to the interests it affects, it is at the point at which animals that have been raised humanely are about to be painlessly killed that the most important question arises – namely, whether the killing can be justified by reference to the interests that are at stake. I have argued that in general it cannot. The animals' interest in continuing to live outweighs the human interest in eating them. That those who now want to kill the animals had earlier caused them to exist – an act that was good for them – is, at this point, irrelevant. One cannot plausibly claim that in killing them one would be depriving them only of what one gave them in the first place. That justification would allow parents to kill their children. Whatever good the practice has bestowed on animals up to this point cannot be cited as credit from which the killing can now be debited.

The argument for having a practice of benign carnivorousness appeals to two considerations: the human interest in eating meat, and whatever *impersonal* reasons one might have to cause animals to exist with lives that would be good for them. In general, we assign little or no weight to impersonal reasons to cause individuals to exist. We do not, for example, accept that there is a significant moral reason to cause a new person to exist simply on the ground that the person's life would

be good.<sup>4</sup> It would be surprising if we thought there were *any* impersonal reason to cause animals to exist simply on the ground that their lives would be good.

(There is, however, a deep, unresolved problem here. Although we deny that there is a significant impersonal reason to cause individuals to exist because their lives would be good, we accept that there is a significant impersonal reason not to cause individuals to exist if their lives would be bad. These intuitions are entirely compelling: while there is no moral pressure to have children, or to breed animals, *just* because they would be happy, there *is* strong moral pressure not to cause people to exist if their lives would be utterly miserable. To my knowledge, no one has offered a satisfactory explanation of this puzzling asymmetry.)

The defender of benign carnivorousness might concede that while there is no strong positive case in favor of the practice, such a case is unnecessary. All that is necessary is that the practice be permissible. Our interest in having it will then supply the motivation to implement it. Yet considerations of interests suggest that it is in fact *not* permissible. Given the interests at stake, we cannot justify the killing that is involved in benign carnivorousness.

Two lines of argument are open to the proponent of benign carnivorousness at this point. First, suppose we have caused certain animals to exist and raised them humanely in order to eat them. We have reached the point at which we planned

<sup>4</sup> That is our intuition in current conditions. But this intuition may reflect a deeper belief that good lives have a diminishing marginal impersonal value. If the human race were on the verge of extinction, we would have a very strong reason to cause new people to exist.

to kill and eat them, but now realize that their interest in continuing to live outweighs our interest in eating them. What is the alternative to killing them? If we now refrain from killing them, are we morally required to continue feeding and caring for them until they die naturally?

If we are *not* required to continue to provide for them, it seems that we must be permitted to release them into the wild. But animals that are bred for human consumption are, like domesticated pets, largely incapable of surviving in the wild. Even the most hardened animal-rights activists usually favor the painless killing of domesticated animals for whom no home can be found. They regard it as a form of euthanasia, since animals unsuited to life in the wild are likely to suffer from hunger and disease before being painfully killed by a predator or an automobile. But if it is better for domesticated animals to be painlessly killed than to be allowed to suffer a slow and miserable death in the wild, it seems permissible after all to kill animals raised as part of the practice of benign carnivorousness. But if we can permissibly kill them, why can we not eat them once they are dead?

What is questionable here is the assumption that one can cause an individual to exist for purposes of one's own without acquiring responsibilities. To cause an individual to exist in a vulnerable and dependent condition is arguably to make oneself liable to certain duties of care. It seems wrong to cause an individual that is incapable of surviving in the wild to exist and then to abandon it in the wild. One must either refrain from causing it to exist or else arrange for it to have the care it requires once it exists.

The second line of argument open to the defender of benign carnivorousness in-

volves distinguishing between the practice as a whole and the act of killing in particular. One can argue that, while killing the animals is bad for them, and worse for them than enabling them to continue to live, the practice as a whole, which includes the act of killing, is good for them. It seems a mistake to allow the evaluation of one component of the practice to determine the value of the practice as a whole. Perhaps we should regard the practice as a whole as the appropriate unit of moral evaluation, and consider the act of killing only insofar as it is a component of the practice.

Debates about both punishment and nuclear deterrence have familiarized us with the idea that the rationality or morality of an act can be determined by the rationality or morality of a strategy or policy in which it is embedded. Some philosophers have argued that if it is permissible to threaten a potential criminal with punishment by programming a device that will automatically punish him if he commits a crime, then it must also be permissible to disaggregate the automatic punishment strategy into its constituent parts by separately threatening punishment and then fulfilling the threat if it is defied. The permissibility of each component is thought to follow from the permissibility of the strategy as a whole.

But this reasoning is mistaken. It can be permissible to bring about a series of effects through a single act, and yet not be permissible to bring about each of the effects through a series of acts. This becomes clear when we consider a parallel argument about nuclear deterrence. Suppose that we could permissibly program an automatic nuclear retaliatory device to annihilate an enemy country if it strikes us first, provided that programming the device would have a high probability of deterring a nuclear first strike that would otherwise be highly proba-

ble. We could also permissibly *threaten* a country with retaliatory annihilation to deter a nuclear first strike. But if this threat were to fail and the enemy country were to launch a first strike, it could not possibly be permissible at that point to fulfill our threat by annihilating the enemy country when doing so would serve no purpose whatsoever. This shows that the permissibility of individual acts is determined by the considerations that favor them at the time of action and cannot be derived from the desirability of the larger practices in which they are embedded.

My rejection of this defense of benign carnivorousness suggests, perhaps surprisingly, that a different form of benign carnivorousness could be permissible. The argument for punishment cited above begins with an example of a single act – the programming of an automatic punishment device – that has two effects: strengthening the deterrence of offenses and imposing a risk of retaliatory harm. The legitimate deterrent aim may justify the risk, thereby making the single act permissible, *even* when it results in the actual infliction of harm.

The problem with the argument is that it does not follow that if each effect were the result of a different act, both acts would be justified. Just as our actual practice of punishment involves two distinct acts – threatening punishment and inflicting it – so benign carnivorousness, as conceived by its proponents, involves both causing animals to exist and then later causing them to cease to exist. But what if we could bundle both these effects into a single act, in the way that making a threat and fulfilling it are bundled together in the programming of the automatic punishment device?

Here is how it might work. Suppose that we could create a breed of animals

*Eating animals the nice way*



genetically programmed to die at a comparatively early age, when their meat would taste best. We could then have a practice of benign carnivorism that would involve causing such animals to exist, raising them for a certain period in conditions in which they would be content, and then simply collecting their bodies for human consumption once they died. Such a practice would not be bad for the animals and would arguably be good for them, since they would have lives worth living and would not have existed at all if not for the practice. And the practice would not involve doing anything to them that would be against their interests, such as killing them.

Note that the practice would not cause the animals to live shorter lives than they might otherwise have had. *Other* animals with a different genetic nature might have been caused to exist instead, and these animals might have lived longer. But none of the animals caused to exist by the practice could have lived longer than they did (unless we also had an antidote to the genetic modification – but for the sake of argument, let us assume that we would not).

This form of benign carnivorism escapes the objection I pressed against the more realistic form that many people have advocated. Yet notice that again a parallel practice involving persons would not be permissible. Again imagine that we have adopted a rigid policy of forcing the birth rate to track the death rate. But we are now considering bringing a limited number of people into existence above the replacement level, but only to use their organs to solve the problem of organ shortages. In this version of the example, however, they would not have to be killed on reaching the age of fifty. They would instead be genetically programmed to die with healthy organs at that age.

I doubt that anyone would find this proposal attractive. And it is not obvious that we could explain the difference between this practice and the parallel form of benign carnivorism by reference to people's rights. For the objection to causing such people to exist does not seem to be that it would violate their rights. Although some defenders of rights might disagree, it would not be wrong to have such a child, when *any* child one might have could inherit a genetic defect that would prevent him from living beyond the age of fifty. Nor would having such a child be permissible only because the procreative rights of the parents would override the rights of the child. Rather, there does not seem to be any right to a possibility of living beyond fifty. So the objection to causing people to exist who would be pre-programmed to die at age fifty must, it seems, be impersonal and comparative in character. That is, it seems wrong to cause such people to exist only because we could cause other people to exist instead who would not have the genetic limitation, despite the fact that causing these *different* people to exist would not address the problem of organ shortages.

One might argue that the objection to this parallel practice involving human beings cannot be simply that it would have been better to cause other people who could have lived longer to exist instead. In the circumstances, it would in fact be *worse* overall to cause such people to exist, since their existence would exacerbate the population problem without solving the organ shortage problem. The real objection, one might argue, concerns equality. The genetically preprogrammed people we might cause to exist would be our moral equals, but we would have deliberately ensured that their lifelong well-being would be lower than that of most other people. To create

a distinct group of people with reduced longevity would be inequalitarian.

Note that this objection also takes an impersonal form. If the inequality created by causing the new people to exist is objectionable, it is not because it is worse, or bad, for the worse-off people. Because these people's lives would be well worth living, it is, if anything, good for them to exist. Inequality that is not worse for anyone may well be morally objectionable, but it is not objectionable enough in this case to explain our sense that it would be wrong to cause these people to exist with a genetically predetermined limit to their longevity.

One other possible explanation is that to cause these people to exist would be to *use* them for the sake of others. Yet that objection may not apply if our policy was never to use such people's organs without their freely given consent. Some might refuse. But we could create just enough for it to be statistically predictable that there would be enough volunteers to solve the problem of organ shortages.

In the hypothetical example, it seems wrong, intuitively, to cause people to exist with a genetically determined maximum life span of fifty years. But it is not clear why exactly this is wrong. This leaves open the possibility that the explanation of why it is wrong, whatever it may be, will apply as well to the second form of benign carnivorism.

Yet there seems to be an interesting difference between causing human beings to exist who are preprogrammed to die prematurely and causing animals to exist that are preprogrammed to die in good health at a certain age. Suppose that in each case the preprogramming was a result of a random mutation rather than of human choice. If some human beings were found to have a gene that caused them to die at age fifty no matter

their state of health, most people would support efforts to eliminate this gene via voluntary selection. That is, most people would favor making it possible for potential parents to have themselves or their embryos screened for the gene in order to prevent the birth of people who would have it. Certainly we would not welcome the presence of this gene because it would help make more organs available for transplantation.

But if we found a naturally occurring strain in some animal species whose members were genetically determined to die prior to the onset of age-related deterioration, we might welcome this discovery as making possible a practice of benign carnivorism that would not require either the killing or the genetic modification of the animals we would consume. There would, it seems, be no more reason to eliminate the gene than there is to try to increase the life spans of shorter-lived species to match those of longer-lived species. This, at any rate, is the common intuition. Whether it is to be trusted is another matter.

The only form of benign carnivorism that is possible now – raising animals humanely and killing them painlessly – seems morally unjustifiable because the interest the animals would have in not being killed would decisively outweigh the interest people would have in killing and eating them. It does not, however, seem morally objectionable to eat an animal that has died of natural causes, which suggests that it could be permissible to use techniques of genetic modification, when they become available, to create animals that would die naturally on a predictable schedule and in good health. It is hard to see what could be wrong with this practice, though a parallel practice involving human beings would not be permissible, which casts

*Eating animals the nice way*

some doubt on the permissibility of the practice involving animals.

We might go further and imagine a version of benign carnivorism based on genetically modifying animals so that they would not only die in a healthy state on a predictable schedule but also enjoy longer lives than their unmodified counterparts. This possibility, however, highlights a problem that afflicts all the variants of benign carnivorism we have considered – namely, that because the animals would be raised in humane conditions and would live for more than just a short period, we would have to invest more in each animal than we currently do in factory-farmed and intensively reared animals. This greater investment would force the unit price up and cause economies of scale to decline. Meat would become a luxury available on a regular basis only to the rich. While this outcome would be objectionable on grounds of equality, it might not be so bad on balance, since decreased consumption of meat would very likely improve the health and longevity of the general population. Almost any shift away from the ways in which meat is currently produced and consumed would be better for both animals and people.<sup>5</sup>

5 I am grateful to Joshua Knobe for stimulating conversation, and to Derek Parfit for extensive and illuminating written comments on an earlier draft of this article.