

Seven Animals and Killing

In setting out to write this paper, my intention was to fill a gap in my book *Animal Liberation*.¹ There I argued that the interests of animals ought to be considered equally with our own interests and that from this equality it follows that we ought to become vegetarian. The argument for vegetarianism is not based on any claim about the wrongness of killing animals – although some careless reviewers read this claim into my book, no doubt because they assumed that any moral argument for vegetarianism *must* be based on the wrongness of killing. Instead the argument for vegetarianism is based on the suffering that is, and as far as I can see always will be, associated with the rearing and slaughtering of animals on a large scale to feed urban populations. I explicitly avoided taking a position on the wrongness of killing animals, for I wanted the book to reach nonphilosophers, and the issue of killing cannot be dealt with briefly and simply.

So why is it wrong to kill? Since my basic ethical position is utilitarian, my initial approach is to ask what objections a utilitarian would have to killing. If we are considering killing a normal human being, the utilitarian can point to obvious bad effects that the killing of one normal human has on others. Killing leads to grief on the part of friends and relatives of the victim, and to fear and insecurity in the community generally. Nonutilitarians, however, regard these as mere “side-effects,” not touching upon the real wrongness of killing: for, they say, it would be wrong to kill even if one’s victim was a hermit whose death would never be discovered. I am not sure that we should, in the case of

normal human beings, allow these “side-effects” to be so lightly brushed aside – the use of fantastic examples can be misleading. In the case of at least some species of nonhuman animals, however, these side-effects cannot explain the wrongness of killing – if it is wrong – since the animals may lack sufficient knowledge of what is happening to feel fear and insecurity, and pair-bonding, maternal relations, or social relations may not be strong enough to give rise to any sense of loss among survivors. (I said in the case of *some* species – with others this may not be true.) This lack of knowledge, incidentally, will also be true of some human beings, namely infants and severe mental defectives.

So what should we say about the wrongness of killing beings to whom the “side-effects” argument does not apply? Here the most obvious answer for the utilitarian to give is that, provided the being is capable of pleasant experiences, to kill it is to reduce the amount of pleasure in the world. Since pleasure is good, this is, other things being equal, wrong. Of course, a similar argument about pain points in the opposite direction, and it is only when we believe that the pleasure a being is likely to experience outweighs the pain it is likely to suffer, that this argument counts against killing. So what this amounts to is that we should not cut short a pleasant life.

This seems simple enough: we value pleasure, killing those who lead pleasant lives eliminates the pleasure they would otherwise experience, therefore such killing is wrong. But stating the argument in this way conceals something which, once noticed, makes the issue anything but simple. There are two ways of reducing the amount of pleasure in the world: one is to eliminate pleasure from the lives of those leading pleasant lives; the other is to eliminate those leading pleasant lives. The former leaves behind beings who experience less pleasure than they otherwise would have. The latter does not. This means that we cannot move automatically from a preference for a pleasant life rather than an unpleasant one, to a preference for a pleasant life rather than no life at all. For, it might be objected, being killed does not make us worse off; it makes us cease to exist. Once we have ceased to exist, we shall not miss the pleasure we would have experienced.

Perhaps this seems sophistical. Well, then, consider the opposite case: a case not of reducing pleasure, but of increasing it. There are two ways of increasing the amount of pleasure in the world: one is to increase the pleasure of those who now exist; the other is to increase the number of those who will lead pleasant lives. If killing those leading pleasant lives is bad because of the loss of pleasure, then it would seem to be good to increase the number of those leading pleasant lives. We could do this by having more children, provided we could reasonably expect their lives to be pleasant, or by rearing large numbers of animals under conditions which would ensure that their lives would be pleasant. But would it really be good to create more pleasure by creating more pleased beings?

First published in *Inquiry*, 22 (1979), pp. 145–56. Reprinted by permission of Taylor and Francis, Oslo, Norway.

This perplexing issue was first raised by Henry Sidgwick and has since been revived by Jan Narveson and Derek Parfit.² There are at least three possible approaches. The first is simply to accept that it is good to increase the amount of pleasure in the world by increasing the number of pleased beings, and bad to reduce the amount of pleasure in the world by reducing the number of pleased beings. This approach, which Sidgwick favored, has the advantage of being straightforward and clearly consistent, but it requires us to hold that if we could increase the number of beings leading pleasant lives without making others worse off, it would be good to do so. To see whether you are troubled by this conclusion, it may be helpful to consider a specific case. Imagine that a couple are trying to decide whether to have children. Suppose that so far as their own happiness is concerned, the advantages and disadvantages balance out. Children will interfere with their careers at a crucial stage of their professional lives, and they will have to give up their favourite recreation, cross-country skiing, for a few years at least. On the other hand they know that, like most parents, they will get joy and fulfillment from having children and watching them develop. Suppose that if others will be affected, the good and bad effects will cancel each other out. Finally, suppose that since the couple could provide their children with a good start in life, it is probable that their children will lead pleasant lives. Should the couple count the likely future pleasure of their children as a significant reason for having children? If we accept this first approach, they should.

This approach is known as the “total” view since on this view we ought to increase the total surplus of pleasure over pain, irrespective of whether this is done by increasing the pleasure of existing beings, or increasing the number of beings who exist.

The alternative approach mentioned by Sidgwick is to aim at the highest possible *average* level of happiness. Sidgwick’s assumption that this is the major alternative to the total view is still commonly made; but the “average” view is most implausible. It implies, as Richard Henson has pointed out, that other things being equal, it would be good to kill all those below the average level of happiness, since this would raise the average – but raising the average gives us a new group of people below the new average, who then also become eligible for elimination – and so on.³

If this is not enough there are other objections to the average view. On this view it would be wrong to bring into a world of extremely happy beings an additional being who would have a happy life, but not quite so happy as the already existing people; but why should this be wrong if all it does is create an additional happy life, making no one worse off? Similarly, on the average view, if the world consists only of utterly miserable beings, it would be good to bring into the world more beings who, though still very miserable, would not be quite so miserable as the already existing beings.⁴

For these reasons I shall not consider the average view any further.

A third approach, and a more plausible alternative to the total view, is to count only beings who already exist, prior to the decision we are taking, or who will exist independently of that decision. This is sometimes called a “person-affecting” view, but this is hardly an appropriate label when we are applying it to animals. I shall refer to it as the “prior existence” view. On this view there is no value in increasing pleasure by creating additional beings. This accords with the intuitive judgments most people seem to make about whether couples ought to have children because of the pleasant lives the children are likely to lead (other things being equal). But how do we square this view with our intuitions about the reverse case, when a couple are considering having a child who, perhaps because it will inherit a genetic defect, would lead a thoroughly miserable life and die before its second birthday? We would think it wrong for a couple knowingly to conceive such a child; but if the pleasure a possible child will experience is not a reason *for* bringing it into the world, why is the pain a possible child will experience a reason *against* bringing it into the world? A convincing explanation of this asymmetry has not, to my knowledge, been produced.⁵

It should now be apparent why I avoided the issue of killing in *Animal Liberation*. The issue forces us to choose between three possible versions of utilitarianism. One can be disregarded, but which of the other two should we choose? The difference is important. If we take the “prior existence” view we shall hold that it is wrong to kill any being whose life is likely to contain, or can be brought to contain, more pleasure than pain. This implies that it is normally wrong to kill animals for food, since we could bring it about that these animals had a few pleasant months or even years before they died – and the brief pleasure we get from eating them would not outweigh this.

The other view – the “total” view – can justify meat-eating. Leslie Stephen implicitly invoked it when he wrote:

Of all the arguments for Vegetarianism none is so weak as the argument from humanity. The pig has a stronger interest than anyone in the demand for bacon. If all the world were Jewish, there would be no pigs at all.⁶

Stephen views animals as if they were replaceable, and with this those who accept the total view must agree. The total version of utilitarianism regards sentient beings as valuable only insofar as they make possible the existence of intrinsically valuable experiences like pleasure. It is as if sentient beings are receptacles of something valuable and it does not matter if a receptacle gets broken, so long as there is another receptacle to which the contents can be transferred without any getting spilt. Although meat-eaters are responsible for

the death of the animal they eat and for the loss of pleasure experienced by that animal, they are also responsible for the creation of more animals, since if no one ate meat there would be no more animals bred for fattening. The loss meat-eaters inflict on one animal is thus balanced, on the total view, by the benefit they confer on the next. We may call this the “replaceability” argument.⁷

The first point to note about the replaceability argument is that even if it is valid when the animals in question have a pleasant life, it would not justify eating the flesh of animals reared in modern “factory farms,” where the animals are so crowded together and restricted in their movements that their lives seem to be more of a burden than a benefit to them.⁸

A second point is that if it is good to create life, then presumably it is good for there to be as many people on our planet as it can possibly hold. With the exception of areas suitable only for pasture, the surface of our globe can support more people if we grow plant foods than if we raise animals.⁹

These two points greatly weaken the replaceability argument as a defense of meat-eating, but they do not go to the heart of the matter. Are sentient beings really replaceable?

Henry Salt, a nineteenth-century English vegetarian and author of a book called *Animals’ Rights*, thought that the argument rested on a simple philosophical error:

The fallacy lies in the confusion of thought which attempts to compare existence with non-existence. A person who is already in existence may feel that he would rather have lived than not, but he must first have the *terra firma* of existence to argue from: the moment he begins to argue as if from the abyss of the non-existent, he talks nonsense, by predicating good or evil, happiness or unhappiness, of that of which we can predicate nothing.¹⁰

When I wrote *Animal Liberation* I accepted Salt’s view. I thought it absurd to talk as if one conferred a favor on a being by bringing it into existence, since at the time one confers this favor there is no being at all.¹¹ But now I am less confident, for three reasons.

First, as we have seen, we do seem to do something bad if we bring a miserable being into existence, and if this is so it is difficult to explain why we do not do something good when we bring a happy being into existence. At the time of writing *Animal Liberation* I thought I could overcome this and other difficulties in the way of an acceptable formulation of the prior-existence view. Derek Parfit has convinced me that the difficulties are more formidable than I had supposed.¹²

Secondly, although it would be wrong to bring into existence a being who will be thoroughly miserable, it does not seem wrong for the government of an

underpopulated country to encourage its people to have more children so that the population will rise by, say, one million. Yet of this million, we can be sure that at least one will be thoroughly miserable. If it is not wrong to create the million, but would be wrong to create the single miserable being, the obvious explanation is that there is value in the creation of the 999,999 – or however many it will be – whose lives are happy.¹³

Thirdly, Salt and other advocates of the prior-existence view need to say something about the point at which a being comes into existence, for this, on their view, marks a morally crucial dividing line. Once a being is in existence, we must maximize its happiness; before it is in existence, there is supposed to be nothing at all to take into account. There is something puzzling about attributing such great moral significance to the moment of birth for, as opponents of abortion have often said, birth does not really make a crucial difference to the fetus. A premature newborn may be less developed than a fetus just before the normal time of delivery. On the other hand – as supporters of abortion have pointed out – it is equally puzzling to make conception mark the dividing line, since the zygote immediately after conception does not seem to have any morally relevant properties not possessed by the egg and sperm before they unite. Why should killing it be worse than using a contraceptive to prevent the egg and sperm uniting?¹⁴

It has been suggested that the development of consciousness, or the capacity to feel, is the essential criterion, but while I accept that the possession of consciousness makes it wrong to cause the conscious being to suffer, or to make its conscious states less pleasurable than they otherwise would be, it is not clear why mere consciousness should be crucial to the wrongness of killing.¹⁵ The search for a morally crucial dividing line leads me to a conclusion very different in its implications from Salt’s. Is it possible that, as Michael Tooley has suggested, the important distinction so far as killing is concerned is the distinction between beings that are merely conscious and those that are self-conscious, in the sense of being able to conceive of themselves as distinct entities, existing over time with a past and a future?¹⁶ If we think of a living creature as a self-conscious individual, leading its own life and with a desire to go on living, the replaceability argument holds little appeal. Salt may be thinking of such beings himself, for he concludes his essay by claiming that Lucretius long ago refuted Stephen’s “vulgar sophism” in the following passage of *De Rerum Natura*:

What loss were ours, if we had known not birth?
Let living men to longer life aspire,
While fond affection binds their hearts to earth:
But who never hath tasted life’s desire,
Unborn, impersonal, can feel no death.

This passage supports the claim that there is a difference between killing living beings who “to longer life aspire” and failing to create a being which, “unborn, impersonal,” can feel no loss of life. But what of a being which, though alive, cannot aspire to longer life because it lacks the conception of itself as a being with a future? This kind of being is, in a sense, impersonal. Perhaps, therefore, in killing it one does it no personal wrong, although one does reduce the quantity of happiness in the universe. This wrong, however, can be counterbalanced by bringing into existence a similar being which will lead an equally happy life.

Classical hedonistic utilitarianism does not support this distinction between personal and impersonal life, but a different variety of utilitarianism, preference utilitarianism, does. This is the form of utilitarianism which Hare claims is implied by the universalizability of our moral concepts.¹⁷ Whether or not one accepts this claim, it is clear that universalizability supports the distinction between self-conscious and merely conscious beings. If I imagine myself in turn as a self-conscious and a merely conscious being, it is only in the former case that I could have a desire to live which will not be fulfilled if I am killed. Hence it is only in the former case that my death is not balanced by the creation of a being with similar prospects of pleasurable experiences. Preference utilitarianism reflects this by taking into account the preferences of all affected by an action, and weighing them according to the strength of preference, under certain conditions of knowledge and reflection. Preference utilitarians count the killing of a being with a preference for continued life as worse than the killing of a being without any such preference. Self-conscious beings therefore are not mere receptacles for containing a certain quantity of pleasure, and are not replaceable.

To take the view that non-self-conscious beings are replaceable is not to say that their interests do not count. I have elsewhere argued that their interests do count. As long as a sentient being is conscious, it has an interest in experiencing as much pleasure and as little pain as possible.¹⁸ Sentience suffices to place a being within the sphere of equal consideration of interests; but it does not mean that the being has a personal interest in continuing to live. For a non-self-conscious being, death is the cessation of experiences, in much the same way that birth is the beginning of experiences. Death cannot be contrary to a preference for continued life, any more than birth could be in accordance with a preference for commencing life. To this extent, with non-self-conscious life, birth and death cancel each other out; whereas with self-conscious beings the fact that once self-conscious one may desire to continue living means that death inflicts a loss for which the birth of another is insufficient gain.

This suggests a possible compromise between the two versions of utilitarianism. We might grant that the total view applies when we are dealing with

beings that do not exist as individuals living their own lives. Here it is appropriate to consider only the total amount of happiness. When we switch our attention to self-conscious beings, however, there is more at stake than impersonal quantities of happiness. We are therefore justified in giving priority to self-conscious beings who exist now or, independently of our decisions, will exist at some future time, rather than to the creation of possible extra beings.

This is, of course, only a suggestion towards a utilitarian answer to the problem of killing. It is not an adequately worked-out answer, because it says nothing about the nature of the priority we give to self-conscious beings. Should it be a lexical priority, so that no loss to a self-conscious existing being is justified by the creation of any number of additional happy beings? Or is there a nonarbitrary way of trading off losses to self-conscious beings against the creation of extra happy beings? These questions need to be answered, but I am not going to attempt to answer them now. Instead I shall return to the practical issue from which I began. Even the incomplete answer to the problem of killing allows us to reach some practical conclusions about the killing of animals and humans.

So far as animals are concerned, it obviously becomes important to try to decide which animals are self-conscious, in the sense of being capable of desiring to go on living. Some philosophers have argued that only a language-user can be self-conscious. I do not find these arguments convincing, but the issue is too large to be considered here. In any case I would be prepared to concede that some of the animals commonly killed for food are not self-conscious – chickens could be an example.¹⁹

Given that an animal belongs to a species incapable of self-consciousness, it follows that it is not wrong to rear and kill it for food, provided that it lives a pleasant life and, after being killed, will be replaced by another animal which will lead a similarly pleasant life and would not have existed if the first animal had not been killed. This means that vegetarianism is not obligatory for those who can obtain meat from animals that they know to have been reared in this manner. In practice, I think this exemption will apply only to those who are able to rear their own animals, or have personal knowledge of the conditions under which the animals they eat were raised and killed. For the reasons given in *Animal Liberation*, I doubt if it would apply to commercially reared and slaughtered animals.

I am sure that some will claim that in taking this view of the killing of some nonhuman animals I am myself guilty of “speciesism” – that is, discrimination against beings because they are not members of our own species. My position is not speciesist, because it does not permit the killing of nonhuman beings on the ground that they are not members of our species, but on the ground that they lack the capacity to desire to go on living. The position applies equally to members of our own species who lack the relevant capacity.

This last consequence strikes many as shocking, for it amounts to saying that infants, for instance, are as replaceable as merely conscious animals. (Potential self-consciousness is not enough, for a potentially self-conscious being has never desired to go on living.) In real life we are not likely to want to kill and replace normal babies. Parents who do not wish to keep their infants can and would normally prefer to give them up for adoption. In the case of defective infants, however, replacement could be a desirable option. Suppose that a couple plan to have two children. Their first child is normal, but the second is diagnosed immediately after birth as a severe case of spina bifida. If it lives, the child will grow up paralyzed from the waist down, incontinent, and mentally retarded – though it might, for all that, have a tolerably pleasant existence if it is intensively cared for. Suppose that the couple do not wish to give the child up to an institution, fearing that it may not receive the best care there. Yet they are equally unhappy at the prospect of trying to bring up such a child. They still want two normal children. They feel that with the burden of a handicapped as well as a normal child to bring up, however, they cannot have another child. The replaceability principle would allow them to kill the defective infant and then go ahead with another pregnancy.

Is this conclusion too shocking to accept? Before answering, consider two currently accepted practices which are, in my view, not fundamentally different. The first is the practice of examining the amniotic fluid of pregnant women who have a higher than average risk of giving birth to a defective child. If the examination reveals that the fetus will have a defect such as spina bifida – or even a less serious defect like hemophilia – the woman is offered, and usually accepts, an abortion. She may then get pregnant again and repeat the process until the tests show that she is carrying a normal fetus. In other words the fetus is treated as replaceable. As I have already said, I do not think the moment of birth marks a morally crucial divide.

The second practice is that of allowing defective newborns to die. It is now perfectly standard – and recognized as such by the Department of Health in the United Kingdom – for the more severe cases of spina bifida to receive no medical treatment other than what is necessary to relieve pain. Operations which would enable the infants to live indefinitely are performed only on the less severely affected children. With the others the avowed aim of not operating is that the infant will die as swiftly and painlessly as possible. Because I agree with James Rachels, Jonathan Glover, and others that the distinction between killing and allowing to die is not of intrinsic moral significance, I do not think that this policy differs greatly from direct killing. (Insofar as it does differ, it probably is worse, since it prolongs the ordeal for infant, parents, and hospital staff.)²⁰

Some may object to the two practices I have just described. I think our previous discussion of killing shows that they are justifiable, and that direct

killing of the newborn infant can also be justifiable. So there is no discrimination on the basis of species. The replaceability principle applies, regardless of species, to beings who have never had the capacity to desire continued life.

Notes

- 1 Peter Singer, *Animal Liberation* (New York: New York Review, 1975; London: Jonathan Cape, 1976).
- 2 See H. Sidgwick, *The Methods of Ethics*, 7th edn. (London: Macmillan, 1907), pp. 414–16; Jan Narveson, "Utilitarianism and New Generations," *Mind*, vol. 76 (1967) and "Moral Problems of Population," *The Monist*, vol. 57 (1973) and reprinted in M. Bayles (ed.), *Ethics and Population* (Cambridge, MA: Schenkman, 1976). Derek Parfit's extremely influential contribution was published in Part 4 of his *Reasons and Persons* (Oxford: Oxford University Press, 1984).
- 3 See Richard Henson, "Utilitarianism and the Wrongness of Killing," *Philosophical Review*, vol. 80 (1971). Henson's objection does not apply if the average view is interpreted as directing us to maximize the average amount of happiness *per life lived*, rather than the average level of happiness of those alive at any given instant. Probably Sidgwick had in mind happiness *per life lived*. The objections in the following paragraph, however, apply to both versions of the average view.
- 4 I owe these points to Derek Parfit.
- 5 I attempted to produce one in "A Utilitarian Population Principle" in M. Bayles (ed.), *Ethics and Population*, p. 93; but Parfit's reply ("On Doing the Best for Our Children" in the same volume, pp. 110–11) has convinced me that the attempt failed.
- 6 From *Social Rights and Duties*, quoted by Henry Salt, "The Logic of the Larder," which appeared in Salt's *The Humanities of Diet* (Manchester: The Vegetarian Society, 1914), and has been reprinted in T. Regan and P. Singer (eds.), *Animal Rights and Human Obligations* (Englewood Cliffs, NJ: Prentice-Hall, 1976).
- 7 Jonathan Glover discusses whether human beings are ever replaceable in *Causing Death and Saving Lives* (Harmondsworth: Penguin, 1977), pp. 72–3, 159–60, 163.
- 8 See *Animal Liberation*, ch. 3.
- 9 I owe this point – which assumes that people are better happiness-producers than nonhumans – to Roslind Godlovitch.
- 10 "The Logic of the Larder," in T. Regan and P. Singer (eds.), *Animal Rights*, p. 186.
- 11 *Animal Liberation*, pp. 240–2.
- 12 See the articles cited in note 5 above.
- 13 See R. I. Sikora, "Is it Wrong to Prevent the Existence of Future Generations?" in *Obligations to Future Generations*, edited by B. Barry and R. I. Sikora (Philadelphia: Temple University Press, 1978), pp. 112–66.
- 14 For a similar point see T. G. Roupas, "Abortion and Simple Consciousness," *Philosophy and Public Affairs*, 7 (1978), pp. 154–83 at p. 167.

- 15 The suggestion that simple consciousness is enough has been made by W. S. Pluhar, "The Value of Life," *Journal of Philosophy*, vol. 74 (1977), pp. 159-72.
- 16 "Abortion and Infanticide," *Philosophy and Public Affairs*, vol. 2 (1972); a revised version appears in J. Feinberg (ed.), *The Problem of Abortion* (Belmont, CA: Wadsworth, 1973).
- 17 See R. M. Hare, "Ethical Theory and Utilitarianism," in H. D. Lewis (ed.), *Contemporary British Philosophy*, 4th series (London: Allen & Unwin, 1976); for Hare's own views on the application of this theory to a situation involving taking life, see his "Abortion and the Golden Rule," *Philosophy and Public Affairs*, vol. 4 (1975), pp. 201-22.
- 18 See *Animal Liberation*, ch. 1.
- 19 The fact that an animal struggles when an attempt to capture or kill it is made does not show that it desires to live; all it can show is that the animal in some way perceives the situation as undesirable, and tries to escape from it. On the other hand, it is difficult to establish that an animal does not have a desire to live and even in the case of a chicken there may be sufficient doubt for it to be better to give the chicken the benefit of the doubt.
- 20 See James Rachels, "Active and Passive Euthanasia," *New England Journal of Medicine*, vol. 292 (Jan 9, 1975); and Jonathan Glover, *Causing Death and Saving Lives*, ch. 7. For an opposed view by a doctor involved in the treatment of spina bifida infants see the comments by John Lorber, "Early Results of Selective Treatment of Spina Bifida Cystica," *British Medical Journal*, 4 (1973), esp. p. 204.
- 21 Portions of this paper appear in a book on *Practical Ethics* published by Cambridge University Press in 1979, second edition 1993. It has been read to several university departments, and I have received many helpful comments. Invidious as it is to single out one or two from all those who have helped, it would be worse still not to thank Michael Lockwood and Derek Parfit for their especially valuable assistance.

Eight

To Do or Not to Do?

In 1939 Otto Schmidt was working as a laboratory assistant at a distinguished medical research institute in Germany. He learned, through chance remarks and his own observation, that another unit of the institute was receiving mentally retarded patients from a nearby asylum, and using them as research subjects. The patients were exposed to various poisonous gases, including nerve gas, and then forced to continue walking up and down an inclined ramp. They frequently vomited, and showed other symptoms of illness; but if they stopped, they were beaten with sticks. After a few days, most patients died from the poison gases they had inhaled; the remainder were put to death.

Schmidt was horrified by his discovery. At first he assumed that the scientists carrying out this research were doing so without authority, and that if the authorities were informed, it would be stopped. But his initial attempts to act on this assumption failed when the director of the institute made it clear that he had special permission from the highest levels to carry out this research "in the interests of the German soldier, who may again be exposed to chemical warfare." Schmidt attempted to contact these higher authorities, but he received no response. He also tried to alert the relatives of the patients, but his inquiries revealed that only patients who had no contact with relatives were selected for the experiments.

There was little more, legally, that Schmidt could do, but the experiments were continuing, and he could not simply forget about them. Therefore he decided on the only course of action he could think of that stood a chance of