

## **How to argue for (and against) ethical veganism**

This is the Author's Penultimate Ms. of a paper that will appear in *The Ethics of Food*  
Eds. Anne Barnhill, Tyler Doggett, and Mark Budolfson

**Tristram McPherson**  
Ohio State University  
[dr.tristram@gmail.com](mailto:dr.tristram@gmail.com)

This paper has two goals. The first is to offer a carefully reasoned argument for *ethical veganism*: the view that it is (at least typically) wrong to eat or otherwise use animal products. The second goal is to give you, the reader, some important tools for developing, evaluating, and replying to reasoned arguments for ethical conclusions. I begin by offering you a brief essay, arguing that it is wrong to eat meat. This essay both introduces central elements of my case for veganism, and serves as one helpful model of a short ethics essay. In the remainder of this paper, I use the model essay as a target, to illustrate important strategies for developing objections to ethical arguments. I will also illustrate a range of important ways for the vegan to reply to these objections. You can use the models and skills I illustrate here in your own essays, and in your reasoned evaluation of ethical arguments. I conclude that the arguments and replies offered in this paper add up to a powerful reasoned case for ethical veganism. You can practice the skills I illustrate here to deciding for yourself – in a reasoned way – whether my conclusion is correct.

I begin with the promised model essay:

### **It is Wrong to Eat Meat**

Most of us think that it would be wrong to adopt a puppy from a shelter, in order to take it home and torture it until it dies. However, we do not think it is wrong to eat a steak for dinner. In this essay, I will argue that these views are hard to square with

each other, and that the second view is false: it is wrong to eat meat. My argument has the following structure:

1. It is wrong to make animals suffer
  2. If it is wrong to make animals suffer, then it is wrong to kill animals
  3. If it is wrong to kill animals, then it is wrong to eat meat
- C. It is wrong to eat meat.

This argument is *valid*. This means that the conclusion must be true if all of the premises are true. I will defend each of these premises in turn.

First, why think that it is wrong to make animals suffer? To begin, think about why it is wrong to make another person suffer. Part of the most plausible explanation is that because suffering is awful to experience, it is wrong to inflict suffering. Because an animal's suffering is awful for it, this explanation entails that it is wrong to make an animal suffer.

This premise of my argument assumes that animals *can* suffer, which is mildly controversial. For example, René Descartes suggested that animals are just complicated machines with no inner lives (1991 [1640], 148). However, Descartes' views are scientifically indefensible (see Allen and Trestman 2014, §7.1), so I set them aside.

You might object to my case for my first premise that it is only wrong to make a creature suffer if that creature is an ethical agent: the sort of being who can be morally responsible for its actions. But this is false. It is wrong to make babies suffer, and they are not ethical agents. You might object that it is only wrong to make *human beings* suffer. This is implausible for several reasons. First, think about torturing a baby: what is wrong with this is surely the nature of the suffering inflicted, not the fact that the baby has a human genetic code. Second, imagine a non-human animal with a

miraculous mutation, which has the ability to speak, reason, and feel as much as you or I do. Surely the mere fact that such an animal is not genetically human does not make it okay to torture it (compare Peter Singer's argument against such 'speciesism' in his 1977). And, finally, think again about the case I began this essay with: it is wrong to torture a puppy. But surely the central explanation here is just the same as with a human victim: torture will inflict horrible suffering on the puppy, and it is wrong to inflict such suffering.

Some authors, like Carl Cohen (1986, 867) insist that all suffering is not equal: human suffering is much more ethically important than animal suffering. My argument is compatible with this thesis. I am not arguing that torturing a puppy is just as bad as torturing a human being. I think the latter is typically much worse. My claim is only that making the puppy suffer is wrong, and that the pleasure a human being might take from torturing it does not justify inflicting that suffering.

Next, I argue that if it is wrong to make animals suffer, it is wrong to kill them. Some people find the idea that it is wrong to kill animals much less intuitive than the idea that it is wrong to make them suffer. However, an example shows that this combination of views – that it is wrong to make animals suffer, but not to kill them – is difficult to defend. Suppose that there is a cow that has a disease that will be fatal unless treated by giving the cow a painful medical operation. If the cow would go on to have a long and pleasant life after the operation, performing this operation seems good, not wrong. This shows that an ordinarily wrongful act – inflicting suffering on a cow – can be permissible if it is necessary to save the cow's life. But if *saving* an animal's life can justify inflicting suffering that would otherwise be wrong, it is hard to understand how *taking* that animal's life could be a matter of ethical indifference.

We can bolster this initial argument by combining it with a plausible explanation of why it is wrong to kill animals. One important reason why killing a *person* is typically wrong is that killing typically deprives the victim of an objectively valuable future. That is, killing someone deprives them of the valuable experiences, activities, projects, etc. that they would otherwise have had (compare Marquis 1989, §II; I do not claim, with Marquis, that this is the ‘primary’ thing wrong with killing). This principle applies to animals as well: just as suffering can make an animal’s life go badly, pleasant experiences can make it go well. So, just as with humans, it is plausible that it is (typically) wrong to kill animals because doing so deprives them of a valuable future.

Finally, I argue that if it is wrong to kill animals, it is wrong to eat meat. Killing and eating are, obviously, not the same thing: in our economically specialized society, many meat-eaters never even see the animals they eat alive, let alone make them suffer or kill them. However, this doesn’t mean that eating meat is okay. To see why, consider an analogy.

There is a new restaurant in town: the food is sensational, and the prices are very low. How do they do it? Here’s how: the owner kidnaps world-class chefs, and enslaves them at the restaurant. Suppose that the owner is connected with the mob, and going to the police would just get you killed. Your patronizing the restaurant does not enslave anyone, but it still seems wrong. The explanation for why it is wrong is roughly that by patronizing the restaurant, you would be *complicit* in wrongdoing: you would be benefiting from a wrongful act (enslavement), while economically supporting the wrongdoer (the slaver).

Making animals suffer may be less awful than enslaving another human being. But the same form of explanation applies to eating meat. The raising of animals for food causes those animals a horrifying amount of suffering, and early death (see Mason and Singer 1990 for some of the literally grisly details). If it is wrong to kill animals and to cause them to suffer, then the industry that produces our meat acts wrongly on a massive scale. It is wrong to eat meat because in doing so you are complicit with that massive and systematic wrongdoing.

In this essay I have argued that it is wrong to eat meat. One clarification of this conclusion is in order: like many ethical claims, it should be read as a claim about what is *typically* true. It is typically wrong for you to break all of my fingers, but if doing so is the only way to prevent nuclear catastrophe, break away! Similarly, there may be unusual circumstances in which it is permissible or even required to eat meat. Nonetheless, if my argument is sound, each of us does wrong almost every time we sit down to a meal that contains meat.

I have written “It is Wrong to Eat Meat” as a model short philosophy essay. Unless your professor tells you otherwise, you would do well to emulate several of the stylistic features of this essay:

- The introduction offers a brief clear motive for the question addressed, states the essay’s thesis, and previews the argument to come;
- The argument of the paper is summarized in valid premise/conclusion form;
- The essay does not waste words: every sentence is dedicated to developing the central argument, explaining a concept, introducing an objection or replying to it, or doing other important work. Even the conclusion does important work, introducing a crucial clarification of the argument.

- The essay does not use lengthy quotes from its sources: instead, it cites those sources after stating (in my own words) key claims that I take from them.

The argument of this essay is also an excellent target for reasoned objections. I now discuss how to offer such objections.

First, let's back up a bit and think about the activity that we are engaged in. We are seeking to make and to evaluate reasoned arguments about ethics. For example, the model essay did not just disagree with the claim that it is okay to eat meat; it offered reasons for thinking that claim is incorrect, and it organized those reasons into an argument. Making an argument does not simply aim to persuade your reader. I know, for example, that no reasoned argument is as likely to change eating habits as grisly video footage of life inside the animal factories that produce our meat. If philosophers aimed simply to persuade, we would write clever advertising, rather than carefully argued essays. Instead, my aim as a philosopher is to seek the truth together with my audience, in a way that respects the ability of each person involved to find the truth herself, using her own ability to reason. My aim now is to offer you some tools to enable you to skillfully engage in this sort of respectful argumentation.

For many of you, the conclusion of the model essay is a challenge to your ethical views. You may be tempted to reply to this sort of challenge by simply disagreeing with the conclusion. Resist this temptation: if an author offers you an argument, and you ignore the argument and simply reject their conclusion, it is very difficult to seek the truth together with you. So, when you are presented with an argument, your central question should be: does this argument give me good reason to accept its conclusion? The model argument appears to be *valid*: the truth of its premises would logically ensure the truth of the conclusion. When you object to a valid argument, you should focus on objecting to its premises, not the conclusion. This is because the argument purports to offer you reasons to accept its conclusion, and if you cannot explain why you should

reject those reasons, you aren't providing a compelling reply to the argument.<sup>1</sup> On the other hand, if you can identify a good reason to reject one of the premises of an argument, you have made an important and constructive contribution, by explaining why a reasonable person should not be persuaded by the argument. This is why it is important to learn how to offer reasoned objections to the premises of an argument.

Developing reasoned objections is in part a creative task, and there is no recipe for doing it well. However, there are several useful general strategies for finding good objections. Taking the model essay as a target, I will introduce some of these strategies, and illustrate them with exemplary objections to the model essay. Another important philosophical skill is to assess the import of potential objections. Because of this, when I consider each objection I will discuss whether the objections can be answered, whether it calls for some amendment to the model essay's argument, or whether it constitutes a promising line of objection to the overall strategy of the model argument. The point of carefully exploring objections and replies is to arrive ultimately at the best arguments that can be made on each side of an ethical issue, like the issue of whether it is wrong to eat meat. Because objections should target the premises of an argument (as I have emphasized), I will organize my discussion by focusing on each premise in turn.

### **Premise One: inflicting suffering**

*Premise One* of the model argument says:

1. It is wrong to make animals suffer

---

<sup>1</sup> This paragraph simplifies in several ways. First, in some arguments the premises (even if true) simply fail to support the conclusion, even given a charitable interpretation. Clearly demonstrating that fact can be a powerful way of objecting to such an argument. Second, some philosophers have argued that it can be legitimate to object to certain arguments as a whole, without criticizing either specific premises or the logical structure of the argument. For discussion relevant to our topic, see McPherson 2014 and *forthcoming-a*.

In this section, I consider objections to this principle that are instances of three general strategies for identifying objections: looking to extreme cases, appealing to an obscured distinction, and appealing to a competing ethical principle.

One excellent way to find objections to ethical principles is to look to extreme cases (Hájek *forthcoming-b*, §4; this and Hájek *forthcoming-a* are excellent sources of heuristics for doing philosophy, although they are most suited for somewhat advanced philosophy students). There are several relevant *types* of extreme cases. One type of extreme case involves *raising the stakes*. Suppose, for example, that some generic supervillain will incinerate the earth unless you torture this puppy. It is surely required (and not wrong) to torture the puppy in that case. So it is not *always* wrong to cause animals to suffer.

It is not enough to find an objection: you should also think about how someone sympathetic to the argument that you are objecting to should reply to your objection. In this case, there is a decisive reply to this objection: the conclusion of the model essay already granted that it is only *typically* wrong to eat meat. Because scenarios involving comic-book supervillains are extremely atypical, this is not an effective objection to the argument of the model essay. There is an important lesson here: make sure that you interpret the argument you are objecting to accurately and fairly. Failure to do this is so common it has its own name: the *straw man* fallacy.

A different sort of extreme case is more potent. If we arranged animals on a continuum of cognitive sophistication, we would notice that puppies (which featured in the model argument) are relatively close to us on that continuum. So: what happens to the model argument as we move to animals farther away from us on that continuum? Here is one salient example: oysters and other bivalves lack brains, and so are almost certainly incapable of suffering. Because one cannot make an oyster suffer, it cannot be wrong to eat an oyster for the reasons suggested in the model



essay.<sup>2</sup> Because there is nothing atypical about eating oysters, this case is an important objection to the argument in the model essay.

One important way to reply to an objection is to concede that it requires one to modify one's argument. This objection to the model argument is powerful, and the best reply is thus concessive. I grant that the model argument does not explain why it is wrong to eat oysters, and so I conclude that the conclusion of the model argument should be restricted to apply only to eating animals that can suffer.

This in turn raises a further question: *which* animals, exactly, can suffer? Here there are formidable methodological barriers to investigation (Allen and Trestman 2014, §4). The core problem is that we have no direct access to animals' experiential states, so we must reason about their inner lives on the basis of behavioral, functional, neurobiological and evolutionary considerations. Unsurprisingly, the strongest case for suffering can be made for mammals, where the evolutionary and neurobiological parallels with humans are closest. However, we should not assume that only mammals can experience pain; some have argued that there is evidence for pain experience in all vertebrates (Varner 2003), and in many cases we may simply lack adequate empirical knowledge to be able to assess the issue. Especially hard cases include cephalopods such as squid, which are behaviorally very sophisticated but evolutionarily distant from us. My approach to this issue invokes a modest sort of precautionary principle: Because we are not in a position to be confident about whether birds, fish, and cephalopods can suffer, we are not in a position to know whether we act wrongly when we eat them. Indifference to the possibility that we act wrongly is a vice, and we should avoid eating these animals on that basis.

---

<sup>2</sup> There might, however, be other arguments that count against eating animals that cannot suffer: for example, one could offer environmental objections to how some such animals are raised or harvested, or appeal to the idea that simply being a fellow animal is morally significant.

A second powerful way to find objections to a premise is to identify an important distinction that the argument for that premise ignores. For example, one could argue that Premise One of the model essay becomes less plausible once we make the distinction between *being in pain* and *suffering*. Some philosophers grant that many animals can be in pain. However, they suggest that suffering requires something in addition to being in pain that most non-human animals lack. For example, perhaps it requires a conscious belief: that *I am having this pain* (for discussion, see Akhtar 2011, 496-499). An objector might argue that it is suffering in this sense – and not merely being in pain – that is ethically significant. If this were true, then my argument would at very least be incomplete: I would need to discuss the nature of suffering more carefully, and then explore which animals can experience it.

The best reply to this objection begins by emphasizing that important issue here is not how we should use the word *suffering* (in philosophy you should usually avoid fighting about how to use words). It is rather whether conscious belief (or something like it) is required for pain to be ethically significant. If we are clear on this point, another extreme case shows why this objection fails. The most intense pains tend to fully occupy us: one is unlikely to be thinking anything – let alone *this pain is happening to me* – when in utter agony. But surely it is wrong to inflict utter agony on someone, because of how awful it feels (compare Rachels 2011, 898). This shows that it can be wrong to inflict pain that does not count as suffering in the objector's stipulated sense. If this is true of agonizing pains, it should be true of less intense pains. And if it is true for our own case, it should be true for animals as well. I thus conclude that this objection fails.

A third way to object to an ethical premise is to identify and defend an independently plausible ethical principle that conflicts with it. You may have encountered such a principle in your previous study of ethics, or you might be able to develop one yourself. One example of this strategy is to argue against Premise One of the model argument by appealing to contractualism,

which is one of the most influential contemporary approaches to ethics and political philosophy. The basic idea of contractualism is that moral (or political) principles are principles that reasonable persons would agree to as rules to govern their lives together. So understood, contractualism can seem to cast serious doubt on the ethical significance of animals. As Peter Carruthers notes, according to the contractualist, “Morality is viewed as constructed *by* human beings, in order to facilitate interactions *between* human beings...” (1992, 102, emphasis his). Because it is hard to see how a principle like Premise One would help to facilitate such interactions, contractualism may seem to give us good reasons to reject this premise.

I have two interlocking replies to this objection. First, the most plausible forms of contractualism do not have the implications that the objector claims. Exemplary here is T. M. Scanlon’s extremely influential contractualist ethical theory. Scanlon is careful to argue that his theory can be extended to protect animals (1998, 177-84). Further, Scanlon is clear that we have strong reasons that are not based in the contractual principle, so his view is compatible with the idea that we might have such reasons not to harm animals. Other philosophers have been more ambitious, offering contractualist arguments on behalf of animals (Rowlands 2002, Ch. 3; Talbert 2006).

Of course, there are some contractualist theories that have the implications that Carruthers suggests. But these are controversial views among contractualists, and contractualism itself is only one of a number of controversial and competing general ethical theories. Because of this controversy, however, it is unlikely that we should be confident in the truth of these contractualist theories. Without such confidence, however, it is hard to see how these theories could give us good reasons to reject Premise One. Further, the case of animals is exactly one where these theories appear implausible. Because it is obviously wrong for me to torture puppies just for fun, it counts against a moral theory that implies otherwise. This sort of case is part of a

deep and more general challenge. As Martha Nussbaum (2006) and others have argued, many of our most important moral concerns address the interests of distinctively vulnerable parties (such as children, the severely mentally handicapped, and animals), and not simply the interactions between equally capable adult humans. A contractualism that ignores these interests is indefensible. In light of these considerations, it will be very difficult to mount a compelling case against Premise One of my argument that appeals to contractualism.

In this section I have considered three kinds of objections to Premise One of my argument. An important part of my reply has been concessive, refining the premise that I want to defend: so refined, the thesis states that it is typically wrong to inflict pain on a range of animals, including at least all mammals. So refined, I have suggested that it is very difficult to reasonably reject this premise.

### **Premise Two: killing animals**

*Premise Two* of the model argument says:

2. If it is wrong to make animals suffer, then it is wrong to kill animals

This premise is *conditional*. In order to successfully object to a conditional, one would need to find a reasonable way to accept the *antecedent* (i.e. the first part) of the conditional, while rejecting the *consequent* (i.e. the second part). Here, this would mean granting that it is wrong to make animals suffer, and arguing that it is nonetheless okay to kill them. I will consider three strategies for objecting to this premise. These strategies all target my explanation of why we should accept this premise. This was the idea that the wrongness of killing is well-explained by the fact that killing deprives the victim of a valuable future. The first strategy appeals to a competing explanation, the second strategy objects that my explanation is incomplete, and the third objects that my explanation has a false presupposition.

Just as a promising objection can be based in a competing ethical principle, so we can base an objection in a competing ethical explanation. In arguing for Premise Two, I offered a general explanation of the wrongness of killing: that killing can be wrong because it deprives the victim of a valuable future. One seemingly competing explanation is that killing you would be wrong because it would violate your autonomy. *Violation* here includes two important and separable ideas. First, killing you would interfere with your exercise of your autonomy. You cannot live your life in the way you choose if you are dead. Second, killing you would be a way of failing to respect your autonomy: if I take myself to be licensed to kill you, I take myself to have the right to ride roughshod over your own view of how your life should go.

In certain cases, this explanation of the wrongness of killing may seem markedly superior to the 'valuable future' explanation offered in the model essay. For example, suppose that Alice is near death and in pain, but wishes to continue living. If I inject her with a lethal dose of morphine, I wrongly kill her. (Notice that this is another instance of using an extreme case to make a point.) The best explanation here is that I have wrongly failed to respect her right to autonomously determine whether she continues to live. By contrast, it is not clear that I deprive her of a valuable future at all. Because most non-human animals are not autonomous agents (there may be borderline cases of non-human agency, such as chimpanzees), this competing explanation suggests that Premise Two is false.

To see why this objection is not promising, notice that the autonomy-violation explanation also clearly fails in some cases. Because you are an autonomous agent, I should not force you to go to bed at a certain time, even if it is good for you to do so. By contrast, I act *rightly* when I paternalistically force my three-year-old son to go to bed at an appropriate time. This is because he is not an autonomous agent. It would obviously be very wrong to kill my son, but since he is not an autonomous agent, this cannot be explained in terms of autonomy violation.

You might think that this leaves us at an impasse: we have two candidate explanations of the wrongness of killing (valuable future-deprivation and autonomy-violation) and counter-examples to each. Does this show that both must be bad explanations? No. A better diagnosis is that each of these accounts provides a *typically sufficient but not necessary* explanation of the wrongness of killing. That is: killing can be wrong *either* because it violates autonomy, *or* because it deprives the victim of a valuable future (or both). On this account, one of the reasons why it is uncontroversial that it is wrong to kill an adult human in a range of ordinary circumstances is that there are several different things wrong with such killing. Because Premise Two of the model argument requires only that the valuable future-deprivation explanation is typically sufficient, this reply vindicates that premise.

Another important way to object to an ethical principle or explanation is to argue that it is *incomplete*. Whenever someone offers a principle or explanation, it is always a good idea to ask: is that *all* that is doing the important explanatory work here? Or have I only been given a part of the best ethical principle that applies to this sort of case? For example, one might think that a future's merely *being valuable* is not sufficient to explain why we must not eliminate it. Suppose that my wife and I were debating whether to have another child. If we did, that child would almost certainly have a valuable future. But it seems clear that I do no wrong simply by preventing that future: my having had a vasectomy does not make me akin to a murderer. Michael Tooley (1972) proposes an explanation of this fact: in order for it to be wrong to deprive an entity of a valuable future, that entity needs to have the capacity to care about its own continued existence. Because the child I do not conceive does not currently exist, it cannot care about its future existence, and hence I do not wrong it by preventing its future, on Tooley's view. One might appeal to Tooley's view to argue that it is not wrong to kill most animals because they are not cognitively sophisticated enough to care about their continued existence.

It is not obvious whether some animals can care about their futures in the relevant way. However, I set this aside, and instead focus on arguing against Tooley's explanatory claim (notice that I do so by appealing to extreme cases reasoning). Suppose that artificial intelligence research advances to the point that we are capable of creating intelligent and autonomous androids, capable of almost everything humans are: sophisticated reasoning, love, physical and emotional pain, etc. Suppose this type of android is programmed to be simply incapable of caring about its own continued existence, although it can and typically does care deeply about particular others. In light of this programming, such androids would be predictably prone to certain tragic behavior: they would sometimes lay down their lives to save others from inconsequential harms. It would be obviously wrong to kill such an android – even with its consent – to save yourself from a splinter. The android's inability to care about its future is a rational imperfection, but not one that licenses killing it.

We can square our judgments about the nonexistent child and the android if we suggest that the android is (imagined to be) an ethically significant being that now exists, while the nonexistent child is a merely possible entity. We can then amend our ethical explanation as follows: it is wrong to deprive existing creatures of valuable futures, but it is not wrong to prevent non-existing entities from coming into existence. This explanation implies that it is wrong to kill existing animals, and so supports Premise Two.

A third way to object to an explanation is to argue that it has a false presupposition. That is: it works only by implicitly assuming some false claim. One crucial presupposition of my explanation of the wrongness of killing animals is that if I refrain from killing a cow today, there is a single moral patient – the cow – that will enjoy various pleasant cow experiences in the future. The objector suggests that we have reason to doubt this, if we think carefully about the conditions for the continued existence of a given moral patient.

The objection can be initially motivated by another extreme case: suppose that a mad scientist was able to map the neural structure of our brains, and then *swap* those structures: your brain is ‘wiped’ and then rebuilt in accordance with the map of my brain, so that your body is now the home of beliefs, desires, and ‘memories’ near identical to mine (pre-operation), and vice-versa. Suppose that the mad scientist performs this swap shortly after I injected my own heart with a slow-acting but lethal poison. Arguably, thanks to the mad scientist’s intervention, I will have survived, and succeeded in killing you, rather than myself.

Some philosophers use cases like these to argue that *psychological continuity* is required for personal identity or ethically significant survival. In the case above, the idea is that I survive the operation because the surviving body houses a psychology that is continuous with my pre-operation psychology. This sort of case matters to my argument for the following reason. Suppose that the psychological connections across the life of a cow are not very rich. Then, in killing the cow now, I may be depriving *it* of only an inconsequential amount of valuable future. I will also be preventing a series of future ‘cows’ from coming into existence and enjoying life. But as we saw from the nonexistent child case, it appears not to be wrong to refrain from bringing into existence beings with valuable futures.

This leads us to the central issue: do animals have rich enough psychological connections to underwrite the intuitive thought that a given cow (e.g.) is the same moral patient over time? I am cautiously optimistic that they do, in at least many cases. For example, many animals appear capable of various forms of memory (Allen and Trestman 2014, §7.4). However, as with questions about animal pain and suffering, answers here are likely to vary substantially across species in ways that require careful empirical work to tease out.

It is worth making two further points. First, theories of personal identity – and related claims about the persistence of a given moral patient – are extremely difficult to assess. The view



that psychological continuity is the criterion of ethically significant survival is controversial. And on many competing views – on which organism continuity, or brain continuity can underwrite ethically significant survival, for example – the objection will fail immediately. Second, the precautionary approach to practical ethics that I advocated in the preceding section is again relevant here. This objection certainly reveals deep complexities ignored by the argument of the model essay. However, we should only be content to reject that argument if these complexities lead us to be confident that it is not wrong to kill animals.

In this section I have argued that the appeal to autonomy-violation complements – rather than competes with – the valuable future-deprivation account of the wrongness of killing. I also argued that it can be wrong to kill a being that is incapable of caring about its own continued existence. In discussing this issue, I amended the ethical principle I endorse to claim that it is wrong to deprive an *existing* moral patient of the valuable future that it would otherwise have. And I argued that many animals are probably the same moral patient across time (although I granted that the issues here are quite complex).

### **Premise Three: using animal products**

*Premise Three* of the model argument is the claim that:

3. If it is wrong to kill animals, then it is wrong to eat meat

I argued for this claim by defending another ethical principle: that it is wrong to be complicit in wrongdoing: to benefit from that wrongdoing, and to support it. Here we can ask two questions: should we accept this principle, and does it really support Premise Three? In this section I begin by considering an objection to the idea that this principle supports Premise Three, before considering whether this principle is objectionably incomplete.

It is always wise to consider whether an ethical principle really supports the conclusion it is intended to. Consider a case that illustrates this sort of objection to Premise Three. Suppose Alice is driving carefully on a country road, when a deer jumps in front of her truck without warning. The deer is killed instantly, and Alice moves its carcass to the side of the road and leaves. Zoe, who lives nearby, sees all of this. Zoe knows how to dress a deer carcass and has a taste for venison. She takes the carcass home, dresses it, cooks some, and eats it (compare Bruckner, *forthcoming* for further discussion of cases like this one). In this scenario, Zoe knowingly prepared and ate meat. But in doing so, Zoe is not complicit in any wrongdoing: Alice's killing of the deer was neither malicious nor negligent, so it is hard to see how it could be wrong. This is a case where eating meat is not complicit in wrongful killing, so it is a case where Premise Three fails to hold, even if the principle I offer is true.

I am happy to grant the objector this case. Recall that my conclusion is that eating meat is *typically* wrong. This case helpfully brings out another atypical exception. The objection lacks more general force exactly because in the overwhelming majority cases, the meat that we eat is wrongfully produced (at least if the arguments for Premises One and Two are sound).

We saw in the previous section that a good strategy for finding objections to an explanation is to challenge its completeness. The same is true for ethical principles like my complicity principle. We can challenge the completeness of my principle by arguing that complicity with the wrongful treatment of animals could only be wrong if it tended to make a difference to how much wrongful treatment there was (see Appiah 1986-7 for a version of this view about complicity). I will call this the *efficacy objection*. This objection has significant force: one might wonder what the point of avoiding complicity is, if it makes no difference to how much animal suffering occurs.

Because I take this to be the single most important challenge to the argument of the model paper, I will offer three potentially complementary replies. (Please note that I offer multiple replies to help illustrate the issues here. In general you should focus on developing the single strongest reply to an objection as clearly as you can, rather than offering multiple replies.)

The first reply accepts the objection, and claims that it is wrong to eat meat because doing so *does* tend to make a difference to the amount of mistreatment of animals. This might seem absurd: by the time I buy a chicken at the store (for example) it is already dead. And the idea that every chicken bought will cause another one to be raised, made to suffer, and then killed, is plainly false. Peter Singer (1980, 335-6, and following him, Norcross 2004, Kagan 2011, and Rachels 2011) has replied to this challenge in the following way: There must be some change in demand for chicken that the market would notice. For example, Singer imagines that for every 10,000 vegetarians, there would be one fewer 20,000-bird broiler factory, harming and killing 100,000 chickens a year. He imagines further that if we were just below the threshold – if xxx9,999 people were vegetarian – the last 9,999 vegetarians would save no chickens, because demand for chicken would be just above the threshold that triggers a change in supply. Given these assumptions, and given that we do not know exactly how many other vegetarians there are, someone becoming vegetarian has only a 1/10,000 chance of making any difference to the number of chickens made to suffer and die.

That sounds depressing. But Singer argues that we should pay attention to the other numbers: if one is that 1/10,000, one will save 100,000 chickens a year. In light of this, the *expected* effect of becoming vegetarian is the effect you would have if you make a difference divided by your chance of making that difference; in the example, saving  $100,000/10,000=10$  chickens a year from short but awful lives. Of course, these precise numbers are merely illustrative; Singer grants that we do not know where exactly the thresholds are. But he suggests that the structure of

probable effects will be similar on any reasonable hypothesis about these thresholds. So, according to Singer, while any reduction or increase in one's meat consumption has a tiny chance of making a difference to the amount of wrongful animal suffering and death, the difference you will make if you do make a difference will be correspondingly huge. And this, it might be claimed, is what makes it wrong to eat meat. If Singer's reasoning is sound, it answers the efficacy objection: complicity is wrong in part because it has an ethically significant chance of making an ethically significant difference. While Singer's reply is promising, his argument is somewhat complex, and relies on some controversial assumptions (see Budolfson *ms.* for an important reply). In light of this, I will explore alternative ways of replying to the efficacy objection.

Singer's argument illustrates two important ideas worth keeping in mind in your ethical reasoning. First, sometimes the *expected* effects of your actions are ethically significant, and not just their actual effects. (In this case, the alleged expected effect of being a vegetarian is sparing ten chickens a year from short and awful lives, even if for most vegetarians, there is no actual effect on chicken well-being.) Second, in thinking about the effects of an action, it is sometimes important to step back from focusing on the particular act, and think about how that act fits into overall patterns.

A second response to the efficacy objection appeals to these patterns in another way, by focusing on the ethical significance of what groups of people do together. One advantage of this approach is that it is uncontroversial that meat-eaters as a group *do* make a difference to the amount of animal suffering: if there were no omnivores there would be no factory farms.

I will introduce the key idea with another example. Suppose that there are two small cities, Upstream and Downstream, along the same river. The river is the only available source of water for the households in each city. Each household in each city draws its water from the river as the river comes into the city, and dumps its sewage in the river as it flows out of the city. The

sewage dumped in the river in Upstream flows down the river and pollutes the drinking water drawn from the river in Downstream. As a result, the people in Downstream are constantly getting seriously ill and dying. Suppose that each household in Upstream could, at small cost, bury their sewage instead of dumping it in the river. If everyone in Upstream did this, it would end the health catastrophe in Downstream. However, given the number of other households that are actually polluting, a single person in Upstream burying his sewage would not save anyone in Downstream from illness or death.

It seems plausible that the sewage-dumpers in Upstream *together* wrongfully cause massive amounts of suffering and death in Downstream. Anyone in Upstream who dumps her sewage in the river is thus part of a group that acts wrongly. It is easy to cease to be part of that group, however: one need only bury one's sewage. It seems plausible that one should bury one's sewage in this situation, rather than dump it into the river. We could explain this by appealing to the following ethical principle: if one can avoid being part of a group that together does serious wrong, then one acts wrongly by continuing to be a member of that group. This principle applies neatly to eating meat. Together, the meat-eaters make a tremendous difference: without their demand for meat, no one would cause animals to suffer and die in order to produce it. So the meat-eaters together make vast amounts of wrongful pain and death happen to animals. So, by the ethical principle just proposed, one acts wrongly by continuing to be a member of that group.

This reply answers the efficacy objection by appealing to group efficacy. However, the issue of when exactly it is wrong to remain a part of an ethically objectionable group is very complicated. (Sometimes, for example, it is only by being part of such a group that one can mitigate the bad things the group does.) So I will explore another alternative response to the efficacy objection, which is the one I find most promising.

This response directly rejects the efficacy objection, and defends the claim that complicity with wrongdoing can be a sufficient explanation for wrongdoing, even if it has no expected bad effects. I will defend this response in three ways: by appealing to a plausibly analogous ethical principle, by clarifying the anti-complicity principle, and by appealing to a variant on an earlier case that helps to distinctively motivate it.

The first thing to notice is that there are other plausible ethical principles that require us to act even when our doing so will not make a difference. For example, the duty of *fair play* requires that one not benefit from successful cooperative institutions without making a fair contribution to them; i.e., that one not *freeride* (see e.g. Klosko 2004). Consider, as an example, sneaking onto a public bus without paying the fare.

Second, it may be useful to more precisely state the principle that I endorse (see McPherson *forthcoming-b* for more detailed discussion):

**Anti-Complicity** It is typically wrong to aim to benefit by cooperating with the wrongful elements of others' plans

When introducing a principle, it is often useful to briefly explain each of the elements of that principle. I now do this for Anti-Complicity. My talk of 'plans' here should not be taken to apply only to patterns of explicit reasoning; rather it should include the pattern of goals that explain an individual's or institution's behavior. If my unconscious desire to humiliate my rival explains all of my behavior, humiliating my rival counts as my plan, even if I would never consciously admit this is what I am up to. My talk of 'benefit' should similarly be read in an expansive way: smoking does more harm than good, but if one seeks the enjoyment of a cigarette, one is aiming at benefit in the sense I am interested in. We should understand 'cooperating' in the following way: our plans often call for others to act in certain ways. For example, if I make widgets for sale, my plan includes others' buying those widgets. Of course, it is not crucial that any particular person buys

my widgets. So anyone who buys a widget counts as cooperating with my plan. Finally plans can be disjunctive: someone can plan to read the newspaper, buy some tools at the store, and then use the tools to torture a puppy. The clearly wrongful part here is the puppy torturing. Buying the tools is instrumental to the wrongful behavior, and is arguably wrongful for that reason, and reading the newspaper is not a wrongful part of the plan. It is most clear that we should not cooperate with the wrongful part of the plan.

Anti-Complicity is plausible in part because it can explain the wrongness of certain acts that cannot be explained by either the group or individual efficacy explanations. Return to the example in the model essay: the restaurant that kidnaps and enslaves chefs to make its food. Suppose the restaurant is *demand-insensitive*: it's partly a money-laundering operation, and so it will remain in business even if no one ever patronizes it. This means that neither an individual, nor the whole group of patrons, have any chance of reducing the amount of slavery in the restaurant by refusing to patronize it. Still, it seems wrong to go to the restaurant and enjoy the fruits of the slave chefs' unwilling labors. Anti-Complicity can explain why, while principles that demand that the individual or group make a difference cannot.

If the arguments earlier in the paper are correct, the meat industry has a wrongful plan: to produce meat in a way that involves egregious amounts of pain and early death, and then to sell that meat. They do not, of course, typically sell it directly to consumers. But consumers buying meat is clearly part of their plan: for if consumers do not buy, then wholesalers will not either, and the meat industry's plan would not be economically viable. (This is why meat-industry groups sometimes advertise directly to consumers: to increase consumer-level demand for their goods.) So, in buying meat, one is cooperating with their wrongful plan. And Anti-Complicity suggests that doing so is typically wrong.

Of the three explanations that I have discussed here (individual efficacy, group efficacy, and Anti-Complicity), I prefer the last. However, it is worth emphasizing that, as with explanation of the wrongness of killing, it is not clear that these explanations compete. Rather, if each is sound, they could be complementary explanations of the wrongness of eating meat. This means that the objector has her work cut out for her. For each of the three explanations that I have discussed, she must either debunk the relevant explanatory principle, or argue that the principle does not entail that it is wrong to eat meat. For example, one might insist that individual efficacy is required for wrongdoing in these cases, and then argue against the Singer-style reasoning. I take this to be the most promising way to reject the argument, but to nonetheless be a very difficult task.

### **The implications of the argument: veganism**

In the previous three sections, I have considered several objections to each of the three premises of the model argument, and refined that argument in light of those objections. I now want to consider the broader implications of the argument, as refined. I will begin by considering two objections to the argument that target gaps between the premises and my overall thesis: ethical veganism.

An obvious but important objection at this stage notes that my aim in this paper is to argue for ethical veganism: the view that it is (at least typically) wrong to eat or otherwise use animal products. However, the model argument concludes only that it is wrong to eat *meat*. The model argument itself thus fails to establish ethical veganism. I grant this objection. However, the argument I have developed in this paper naturally extends to support ethical veganism.

The first point to notice is that it is possible to imagine farming with animals in a way that does not involve shortening their lives or making them suffer. My argument does not suggest any



objection to using animal products made on such farms. However, when we turn from possible to actual animal farming, we find that my case against killing animals and making them suffer applies to almost all of the institutions that produce animal products (with the exceptions of some shellfish farms). The reasons lie in the interaction between biology and economics. Consider a single example: even the most humane dairy farm will typically produce as many male calves as female, and almost all of the males will be killed early, so as not to be an economic burden. That means that the central plan of almost any economically viable dairy farming operation involves raising cows to be killed (or to be sold to another operation, knowing the latter operation will kill them), a practice that I have argued above is typically wrong. And this in turn means that the overall argument I have proposed applies here: the core plan of economically viable dairy farms involves systematic wrongdoing, and I have argued that it is wrong to be complicit with such wrongdoing. But one would be complicit with such wrongdoing if one were to buy and consume the milk (e.g.) produced on such farms, and hence buying and consuming such milk would be wrong. This example generalizes to the institutions that produce almost all of our animal products: eggs, cheese, leather, etc. And for this reason I think that my argument supports ethical veganism as opposed to a requirement to be a vegetarian who merely refrains from eating meat.

A second worry about my overall argument is that the initial simple statement of the argument in premise and conclusion form in the model essay is misleading. In the preceding sections, I have emphasized various ways that this argument should be refined, but there is a general worry that should be explored. The conclusion of the model argument emphasized that it is only *typically* wrong to eat meat. And as I explained in my initial discussion of Premise One, this qualifier should be read back into the premises. So the argument should look like this:

1. It is typically wrong to make animals suffer
2. If it is typically wrong to make animals suffer, then it is typically wrong to kill animals

3. If it is typically wrong to kill animals, then it is typically wrong to eat meat

C. It is typically wrong to eat meat.

The first thing to do is to verify that this statement of the argument, like the statement in the model essay, is valid. It is: the addition of the word “typically” does not alter the logical form of the argument, which is: *P, if P then Q, if Q then R, so R*, which is a slightly more complex variant of the classic *modus ponens* argument form. However, there are two connected worries about the argument as given. First, the reference to typicality points us at a range of ordinary cases, but every sort of exception that we have identified for each premise is an exception that must hold for the argument as a whole. The discussion has identified a raft of ‘atypical’ exception cases: cases of making animals suffer to avoid ethically awful alternatives, cases of eating oysters and other animals incapable of experiencing pain, and cases of eating meat (like some roadkill) that was not wrongfully produced.

These cases do not exhaust the set of potential exceptions that the argument permits. And one might worry that as a result, the argument might be far too weak to support anything resembling veganism. To begin to see the force of this worry, notice that I have granted that human suffering and death may tend to be substantially more ethically significant than the suffering and death of non-human animals. This is because, as we saw above, killing you or making you suffer would be wrong for multiple reasons: some have to do with the awfulness of suffering, and the deprivation of your future, and others have to do with the ethical significance of your autonomy.

In light of this, my argument at least suggests that the most central and pressing human interests should typically take priority over the welfare of non-human animals. For example, my conclusion is compatible with the idea that we should typically harm or kill a non-human animal if doing so is needed to prevent suffering or death to a human being. This is practically relevant:

in various times and places, animal products have been an essential element of the only feasible nutritionally adequate human diets. For example, in many parts of the world, owning a cow – or even a handful of chickens – can offer crucial protection against certain forms of malnutrition. I take it to be a virtue of my argument that it is compatible with cases like these counting as legitimate exceptions to the vegan principle.

At this point, however, one may wonder whether the case for veganism has any practical bite at all. After all, becoming a vegan involves a non-trivial sacrifice of real goods. Consider three sorts of examples. First, there is a sea of delicious animal-involving food, so as a vegan one sacrifices access to a range of interesting aesthetic goods. Second, food is deeply meaningful to many people, and animal products are centrally involved in many important cultural traditions and occasions. To be a vegan is thus to complicate one's relationship to those traditions and meanings. Finally, because shared values are central to many personal and professional relationships, veganism could be an impediment to such relationships, especially in cultural contexts where veganism is seen as threatening.

If the fact that veganism required one to sacrifice goods like these typically rendered omnivorism permissible, then veganism would not typically be ethically required, as I claim. Instead, it would be an admirable but non-obligatory ideal. However, I do not think that the sorts of sacrifices just mentioned suffice to make omnivorism permissible. The core issue here is how weighty the considerations in favor of veganism that I have developed in this paper are. The issue is complex, but I think that a reasonable heuristic can be derived from the initial example in the model paper: *some circumstances* would warrant torturing the stray puppy imagined in that example. But those circumstances would be comparatively dire. I contend that only similarly dire circumstances would warrant ordering the sirloin steak for dinner. And the sacrifices typically involved in becoming vegan, while significant, fall well below this threshold.

Challenging this heuristic would be yet another natural way to object to the argument of this paper. However, I think the heuristic is basically sound. And if it is, the argument of the paper suggests that we ought to eschew almost all animal products in almost all ordinary circumstances. We ought, in other words, to be vegan.

### **Pedagogical coda**

I conclude this paper by returning to my pedagogical aims: to aid you in thinking about how to make (and critically examine) philosophical arguments in ethics. In order to do so, I will review the basic elements of philosophical argument that I have sought to explain and illustrate in this paper.

The argument of the model essay began with a vivid *example*: the claim that it would be wrong to torture a puppy in a specific scenario. This case supported a *general principle*: that it is wrong to make animals suffer. The case supported the principle in part because the case seems *representative* of the principle. It does not seem that there is some unique feature of puppies that explains why it is wrong to torture them, for example. This principle was also supported by an underlying *ethical explanation*: it is wrong to cause animals to suffer, *because of how awful it is to experience suffering*. Although general principles themselves appear explanatorily illuminating, explanations and general principles can be different. One way to see this is to notice that there can be multiple good explanations of a single ethical principle, as in the case of the future-deprivation and autonomy-violation explanations of the wrongness of killing. In making arguments, cases, general principles, and explanations are likely to be the most important elements to develop. These elements should fit together in a rationally compelling way, and one good way to do that is to put these elements together into a valid argument, as I again did in the model essay.

In critically examining an argument, you might in principle target the structure of the argument itself: for example, showing that the argument contains some fallacious reasoning. Or you might challenge the ethical claims the argument makes about specific cases: for example, you could try to argue that there is nothing wrong with torturing puppies. (But I dearly hope you don't do that!) It is far more common for it to be useful to challenge the general principles and explanations offered in an argument. I have discussed several important ways of executing these challenges. First, it can be useful to look to extreme case, to see if principles are really generally applicable. For example, the model argument appeals to suffering to explain why it is wrong to eat meat. But some animals (such as oysters) cannot suffer. So the model argument cannot explain why it is wrong to eat them. Second, it can be useful to see if an argument only works because it obscures an ethically important distinction. I discussed this issue using the example of the contrast between being in pain and suffering. Third, it is always a good idea to ask if there is a superior competitor to the general principle or ethical explanation offered in an argument. Examples of this strategy discussed above were the objection from contractualism, and the objection that autonomy violation is the best explanation of why killing is typically wrong. Fourth, a very natural objection to an explanation or principle is that it is *incomplete*. For example, I considered the idea that it is only wrong to deprive a creature of its valuable future *if* that creature is capable of caring about that future. And I discussed the idea that complicity with wrongdoing is only morally objectionable if such complicity can make a difference to the extent of the underlying wrongdoing. Fifth, another important type of objection to some explanations is that those explanations rest on false presuppositions. For example, I considered the possibility that most animals do not have valuable futures in the ethically relevant sense, because most animals lack rich enough psychological connections to remain the same moral patient from one day to the next. Finally, another important way to challenge an argument is to show that there is a gap

between an explanation offered in support of an ethical principle, and the principle itself. For example, the case of eating blamelessly produced roadkill suggests that there is a gap between its being wrong to kill animals, and its being wrong to eat meat, because some meat does not come from animals that were wrongfully killed.

I take these to be some of the most important tools for critically analyzing philosophical arguments. However, there are many more to be discovered. One very good habit to get into when reading philosophical papers is to ask: what kind of argument is this? How is this author objecting to that argument? If you do that consistently, you will soon have a very rich repertoire of tools for evaluating others' arguments, and making your own. One final note about how to use these tools. Probably the most important place to use the tools I have discussed in this paper is in revising your own paper. Once you have a draft of your paper in hand, you should be merciless in carefully reading through it, asking: how compelling is this argument? How could someone reasonably object to it? Are their objections sound? In my view, it is most important to use these tools to examine arguments for the conclusions that you most care about. Only by doing so can you determine whether these conclusions are reasonable, or whether you are guilty of wishful thinking, only accepting them *because* you care about them.<sup>3</sup>

---

<sup>3</sup> One final bit of guidance: you should always recognize help you have received in writing a paper! I am indebted to many people for helpful comments and discussion of ideas related to this paper. These include Mark Budolfson, David Plunkett, Tyler Doggett, Andrew Chignell, Sean Walsh, Derek Baker, Tom Dougherty, Gideon Rosen, and Katie Batterman, to audiences at Rhodes College, Bowling Green State University, Charles Sturt University, and Virginia Commonwealth University, and to many of my students for discussion. I am also indebted to Liz Harman, whose talk about ethical vegetarianism first started me thinking systematically about it. Parts of this paper draw significantly on my 2014, and especially my *forthcoming-b*.

**Works Cited:**

- Akhtar, Sahar. "Animal Pain and Welfare: Can Pain Sometimes Be Worse for Them Than for Us?" *Oxford Handbook of Animal Ethics*. Eds. Tom Beauchamp and R. G. Frey. Oxford University Press, 2011. 495-518.
- Allen, Colin, and Michael Trestman. "Animal Consciousness." *The Stanford Encyclopedia of Philosophy* (Summer 2014 Edition). Ed. Edward N. Zalta.  
<<http://plato.stanford.edu/archives/sum2014/entries/consciousness-animal/>>.
- Bruckner, Donald W. "Why Strict Vegetarianism is Immoral" *Stirring the Pot*. Eds. Robert Fischer and Ben Bramble. Oxford University Press, *forthcoming*.
- Budolfson, Mark. "The Inefficacy Objection and the Problem with the Singer/Norcross/Kagan Response" *Ms.* (cited as *ms-a*).
- Carruthers, Peter. *The Animals Issue*. Cambridge: Cambridge UP, 1992.
- Cohen, Carl. "The Case for the Use of Animals in Biomedical Research." *The New England Journal of Medicine* 315, 1986. 865-869.
- Descartes, Rene. *The Philosophical Writings of Descartes Vol. III*. Ed. and Tr. John Cottingham, Robert Stoothoff, Dugald Murdoch, and Anthony Kenny. Cambridge, Cambridge UP, 1991.
- Hájek, Alan. "Heuristics for Philosophical Creativity." *forthcoming* in *The Philosophy of Creativity*. Eds. Elliot Samuel Paul and Scott Barry Kaufman. Oxford: Oxford University Press. *Cited as forthcoming-a*.
- . "Philosophical Heuristics and Philosophical Methodology." *forthcoming* in *The Oxford Handbook of Philosophical Methodology*. Eds. Herman Cappelen, Tamar Gendler, and John Hawthorne. Oxford: Oxford University Press. *Cited as forthcoming-b*.
- Kagan, Shelly. "Do I Make a Difference?" *Philosophy and Public Affairs* 39.2 (2011). 105-141.

Klosko, George. *The Principle of Fairness and Political Obligation* (New Edition). Rowman and Littlefield, 2004.

McPherson, Tristram. "A Case for Ethical Veganism." *Journal of Moral Philosophy* 11(6), 2014. 677-703.

---. "A Moorean Defense of the Omnivore?" *Forthcoming in The Moral Complexities of Eating Meat*, Eds. Robert Fischer and Ben Bramble. Oxford University Press. Cited as *forthcoming-a*.

---. "Why I am a Vegan (and You Should be One Too)." *Forthcoming in Philosophy Comes to Dinner*, Eds. Andrew Chignell, Terence Cuneo, and Matthew Halterman. Routledge. Cited as *forthcoming-b*.

Marquis, Don. "Why Abortion is Immoral." *Journal of Philosophy* 86.4, April 1989. 183-202.

Mason, Jim and Peter Singer. *Animal Factories* Revised and Updated Edition. Harmony: 1990.

Norcross, Alasdair. "Puppies, Pigs, and People: Eating Meat and Marginal Cases." *Philosophical Perspectives* 18 (Ethics), 2004. 229-245.

Nussbaum, Martha. *Frontiers of Justice*. Harvard: Belknap, 2006.

Rachels, Stuart. "Vegetarianism." *Oxford Handbook of Animal Ethics*. Eds. Tom Beauchamp and R. G. Frey. Oxford: Oxford University Press, 2011. 877-905.

Rowlands, Mark. *Animals Like Us*. Verso: 2002.

Scanlon, T. M. *What We Owe to Each Other*. Cambridge, MA: Belknap: 1998.

Singer, Peter. *Animal Liberation*. New York: Avon, 1977.

---. "Utilitarianism and Vegetarianism." *Philosophy and Public Affairs*. 9.4 (Summer 1980), 325-337.

Talbert, Matthew. "Contractualism and Our Duties to Nonhuman Animals." *Environmental Ethics* 28 (Summer 2006). 202-215.

Tooley, Michael. "Abortion and Infanticide." *Philosophy and Public Affairs* 2.1 (August 1972). 37-65.



Varner, Gary. *In Nature's Interests*. Oxford: Oxford UP, 1998.