

The Moral Problem of Other Minds

Jeff Sebo

Abstract

In this paper I ask how we should treat other beings in cases of uncertainty about sentience. I evaluate three options: (1) an incautious principle that permits us to treat other beings as non-sentient, (2) a precautionary principle that requires us to treat other beings as sentient, and (3) an expected value principle that requires us to multiply our degree of confidence that other beings are sentient by the amount of moral value that they would have if they were. I then draw three conclusions. First, the precautionary and expected value principles are more plausible than the incautious principle. Second, if we accept a precautionary or expected value principle, then we morally ought to treat many beings as having at least partial moral status. Third, if we morally ought to treat many beings as having at least partial moral status, then morality involves more cluelessness and demandingness than we might have thought.

1. Introduction

In 2003 David Foster Wallace took a trip to Maine to write an article for *Gourmet Magazine* about what it was like to attend the Maine Lobster Festival. But the article that he ended up publishing, titled “Consider the Lobster,” is less a travelogue than a deeply personal examination of whether or

not we have a moral duty not to kill lobsters for food.¹ About halfway through this article, Wallace identifies a problem that, I think, deserves more philosophical attention than it receives. He writes:

[T]he questions of whether and how different kinds of animals feel pain ... turn out to be extremely complex and difficult. And comparative neuroanatomy is only part of the problem. ... [The] principles by which we can infer that others experience pain ... involve hard-core philosophy – metaphysics, epistemology, value theory. ... And everything gets progressively more abstract and convoluted as we move farther and farther out from the higher-type mammals into cattle and swine and dogs and cats and rodents, and then birds and fish, and finally invertebrates like lobsters.²

The problem that Wallace is directing our attention to, which I will call *the moral problem of other minds*, works as follows: Suppose that (1) sentience is necessary and sufficient for moral status, (2) we are not always certain if other beings are sentient, and, so, (3) we are not always certain if other beings have moral status. How should we treat other beings in such cases of uncertainty?

Wallace does not try to solve this problem in his article. Instead, he argues that we have some reason to think that lobsters are sentient, and he concludes by expressing confusion about what to do all things considered. The philosophical literature is similar, though less candid. Many philosophers are aware of the moral problem of other minds, but few discuss it in detail, let alone try to solve it.³ Instead, most philosophers try to avoid this problem in one of two ways. First, they argue that particular beings (e.g., fetuses, animals, plants, and sophisticated artificial intelligences) either have

¹ David Foster Wallace, “Consider the Lobster,” *Gourmet Magazine*, August 2004.

² Wallace, “Consider the Lobster,” 5.

³ For an exception, see Guerrero 2007. Guerrero argues that we should not kill individuals unnecessarily in cases of uncertainty about moral status. I think that this principle is compatible with both the precautionary and expected value principles that I will consider here.

or lack sentience.⁴ Second, they stipulate a solution to this problem (e.g., that we should err on the side of caution) without saying much about how this solution works or why we should accept it.⁵

My aim in this paper is to present and evaluate three solutions to the moral problem of other minds. In particular, I will consider (1) an incautious principle that permits us to treat other beings as non-sentient in cases of uncertainty, (2) a precautionary principle that requires us to treat other beings as sentient in cases of uncertainty, and (3) an expected value principle that requires us to multiply our degree of confidence that other beings are sentient by the amount of moral value they would have if they were. I will then argue for three conclusions. First, the precautionary and expected value principles are more plausible than the incautious principle. Second, if we accept the precautionary or expected value principle, then we morally ought to treat many beings, including but not limited to animals, as having at least partial moral status. Third, if we morally ought to treat many beings, including but not limited to animals, as having at least partial moral status, then morality involves more cluelessness and demandingness than we might have thought.

I will proceed as follows. In §2, I will present the assumptions in moral philosophy and philosophy of mind that I will be making here, along with clarifications about the scope of my discussion. In §3-4, I will present and evaluate the incautious, precautionary, and expected value principles. Finally, in §5, I will draw my conclusions and discuss the respects in which they are revisionary.

2. Assumptions and Clarifications

⁴ See, for example, DeGrazia 1996.

⁵ See, for example, Regan 2004, 367.

I begin by presenting the assumptions in moral philosophy and philosophy of mind that I will be making here, along with clarifications about the scope of my discussion.

The first assumption that I will be making here is *sentientism about moral status*. On this assumption, you have moral status, i.e. you are a subject of moral concern, if and only if you are sentient, i.e. if and only if you are capable of phenomenally consciously experiencing pleasure or pain. What does it mean for you to be a subject of moral concern? It means that moral agents can have moral duties to you (for example a duty not to harm you unnecessarily) for your own sake. What does it mean for you to phenomenally consciously experience pleasure or pain? It means that there is something that it is like for you to be experiencing pleasure or pain.⁶ As we will see, this requirement that you have a private, subjective, qualitative experience of pleasure or pain is part of what makes the moral problem of other minds so challenging.

The second assumption that I will be making here is *uncertainty about sentience*. On this assumption, we are not always certain which beings are sentient. This uncertainty has scientific and philosophical sources. Scientifically, we are not always certain which beings experience pleasure and pain. Granted, we can be relatively confident that other vertebrates experience pleasure and pain, since other vertebrates are relatively behaviorally, physiologically, and evolutionarily continuous with humans.⁷ But we cannot be as confident about invertebrates. With respect to some species, like squids, many people think that the continuities with humans outweigh the discontinuities, and so they proceed on the assumption that these animals do experience pleasure and pain. With respect to other species, like ants, many people think that the discontinuities with humans outweigh the continuities, and so they proceed on the assumption that these animals do *not* experience pleasure and pain (a conclusion that, I should note, is less warranted). With respect to still other species, like lobsters,

⁶ For discussion of the distinction between access consciousness and phenomenal consciousness, see Block 1995.

⁷ For discussion of evidence that fish experience pain, see Braithwaite 2010.

many people have no idea what to think. And of course, similar questions arise, to greater and lesser degrees, for other kinds of beings as well, including fetuses, plants, sophisticated artificial intelligences, and even certain kinds of collectives.⁸

Moreover, philosophically, we are not always certain which beings are phenomenally conscious. In particular, many philosophers believe that your private mental life, unlike your behavior and physiology, is not publicly observable even in principle. Granted, I might *perceive* you as having private, subjective experiences, and I might also *infer* that you do by analogy with me. But this perception may be inaccurate, and this inference may be based on bad reasoning, since, after all, we do not generally think that we are justified in drawing an inference about every member of a particular group (e.g., about all humans or animals or living beings) based on our observation of a single member of that group (e.g., ourselves). This is known as the *epistemological problem of other minds*.⁹ Of course, we may think that this problem has a solution, and therefore we may think that we can be relatively confident that some beings are phenomenally conscious. But even if we do think so, it is not likely that this solution will ground certainty about which beings are phenomenally conscious in all cases. Instead, and at most, it will ground near certainty that, say, many other humans are phenomenally conscious, and then a decreasing degree of confidence that other beings are phenomenally conscious depending on the details.

If we combine these assumptions, then they commit us to *uncertainty about moral status*, according to which we are not always certain which beings have moral status. This uncertainty will likely last for a long time, if not forever. Yet in the meantime, we still have to decide how to treat many beings about whom we are uncertain. How many exactly? It is impossible to say. But, to put things in perspective, the Encyclopedia Smithsonian estimates that, at any given time, “there are

⁸ For a discussion of the possibility of collective mentality, see Huebner 2004.

⁹ For discussion of the epistemological problem of other minds, see Carruthers 2004.

some 10 quintillion (10,000,000,000,000,000) individual insects alive.”¹⁰ Moreover, many experts predict that artificial intelligence will soon rival human intelligence, which means that we will soon have to seriously consider the possibility that many artificial beings are sentient too (and they will have to seriously consider the possibility that we are in turn). As a result, it is important that, rather than wait for developments in science and philosophy that may or may not ever arrive, we think as carefully as possible, right now, about how we morally ought to treat others in cases of uncertainty about sentience and moral status. So that is what I will attempt to do here.

Before I do that, however, I think that it will be useful to provide several clarifications about my discussion in what follows. First, this is not a paper about epistemology or philosophy of mind. Instead, it is a paper about ethics. As a result, I will not be asking what we should *believe* about other beings given our evidence about them. Instead, I will be asking how we should *treat* other beings given our beliefs about them. Thus, I will simply stipulate certain beliefs about other minds for the sake of argument (with no expectation that these stipulated beliefs are accurate), so that we can ask what we should do in light of these beliefs. Of course, there are many related questions that we can ask about our beliefs, including what they should be as well as how we should evaluate people whose beliefs are false, irrational, or products of negligence. But I will not be asking those questions here.

Second, this is not a paper about how we should treat other beings *all things considered* in cases of uncertainty about sentience. Instead, this is a paper about whether we should treat other beings *as having moral status* in cases of uncertainty about sentience. As a result, I will not be asking if we can permissibly harm individuals in complex cases involving many conflicting values (since, in these cases, we might think that we can permissibly harm an individual all things considered even if they do have moral status). Instead, I will be asking if we can permissibly harm individuals in simple

¹⁰ “Number of Insects,” Smithsonian Encyclopedia, accessed May 2, 2018, http://www.si.edu/encyclopedia_si/nmnh/buginfo/bugnos.htm

cases involving few if any conflicting values (since, in these cases, we might think that whether we can permissibly harm an individual depends entirely on whether or not they have moral status). Of course, there are many other questions that we need to ask before we can decide what if anything we owe to, say, insects or robots all things considered. But I will not be asking those questions here.

Third, some philosophers think that sentience and moral status can come in degrees. They also think that, if sentient beings have certain other capacities (such as language and rationality) or relationships (such as bonds of care and interdependence), then they can have new kinds of conscious experience and, therefore, new kinds of moral status. On this view, the question is not: Should we treat other beings as sentient and, therefore, as having moral status? The question is rather: Should we treat other beings as having such and such kinds or degrees of sentience and, therefore, should we treat them as having such and such kinds or degrees of moral status (taking into account everything that we know about them)? I think that these complications are important. However, for the sake of simplicity I will focus for the most part on whether or not we should treat other beings as having sentience and, therefore, as having moral status at all. If we like, we can then extend the tools that I will develop for this simple question to other, more complex questions as well.

Finally, and relatedly, the moral problem of other minds arises for other theories of moral status as well. For example, if we accept *rationalism* about moral status, then we will need to ask how to treat particular beings in cases of uncertainty about whether or not they are rational. If we accept *biocentrism* about moral status, then we will need to ask how to treat particular beings in cases of uncertainty about whether or not they are alive. If we accept a *relational approach*, then we will need to ask how to treat particular beings in cases of uncertainty about whether or not they stand in certain kinds of relationships. And so on. And of course, if we are uncertain about which theory of moral status to accept in the first place, then we will have to ask this question at multiple levels, i.e. for every potentially morally relevant capacity or relation, we will have to ask both (1) what to do

given uncertainty about whether this capacity or relation is morally relevant (this is an example of moral uncertainty), and (2) what to do given uncertainty about whether or not a particular individual has this capacity or relation (this is an example of descriptive uncertainty).

My sense is that the discussion will have a similar structure no matter which theory of moral status we accept: The incautious, precautionary, and expected value principles will be our main options in cases of uncertainty, and these options will have similar strengths and limitations. However, I also think that some of the details will depend on which theory of moral status we accept. For example, if we think that we have less evidence about sentience than about rationality, then we might think that cluelessness is more of a problem for sentientism than for rationalism. Similarly, if we think that more beings are possibly sentient than possibly rational, then we might think that demandingness is more of a problem for sentientism than for rationalism. As for uncertainty about which capacities or relations are morally relevant: This kind of moral uncertainty raises special challenges that I will not be able to consider here. Still, I hope that my discussion of the moral problem of other minds will be useful for people independently of which theory of moral status they accept, since this discussion can serve as a blueprint for how other, similar discussions might go.

With these assumptions and clarifications in mind, I will now present and evaluate the incautious principle, the precautionary principle, and the expected value principle.

3. Incautious and precautionary principles

The incautious principle permits treating other beings as non-sentient in cases of uncertainty. In contrast, the precautionary principle requires treating other beings as sentient in cases of uncertainty. As we will see, these principles have symmetrical risks and benefits: The incautious principle risks leading to false negatives (i.e., to mistakenly treating sentient beings as non-sentient), whereas the

precautionary principle risks leading to false positives (i.e., to mistakenly treating non-sentient beings as sentient). There will be other issues as well. However, I will argue that the precautionary principle is more plausible than the incautious principle, since the harm of false negatives is worse than the harm of false positives in this context.

3.1. The incautious principle

Consider first the incautious principle, which permits treating other beings as non-sentient in cases of uncertainty. As far as I can tell, philosophers do not tend to defend incautious principles in other contexts. But I am starting with it anyway since many people seem to accept it in this context.

So, what should we think about the incautious principle in this context? It faces at least two objections. The first objection is that the incautious principle, as currently stated, is too extreme. In particular, it implies that, if you are less than 100% certain that the epistemological problem of other minds has a solution (and, therefore, if you are less than 100% certain that other individuals are phenomenally conscious), then you are morally permitted to treat *everyone in the world other than yourself* as non-sentient. But I think that most of us would agree that this kind of moral solipsism is a non-starter. The mere fact that many other humans are only, say, 99.99999% likely to be sentient, given your evidence, does not justify treating them as though they are not.

A proponent of the incautious principle has three possible replies to this objection. First, they can deny that the incautious principle implies moral solipsism. For example, they can assert that they are, in fact, 100% certain that the epistemological problem of other minds has a solution. In this case, the incautious principle would not imply that they are morally permitted to treat everyone else as non-sentient. Of course, given how hard the epistemological problem of other minds is to solve, this reply strikes me as implausible. But perhaps one can attempt it.

Second, a proponent of the incautious principle can deny that moral solipsism is a non-starter. For example, they can accept the idea that morality is an illusion, and they can try to make this idea seem plausible by noting that, even if morality is an illusion, we might still be motivated to treat others well (e.g., we have self-interested reason to treat others well insofar as doing so benefits us, and we have other-interested reason to treat others well insofar as we care about them). This reply strikes me as implausible too, but, again, perhaps one can attempt it.

Third, a proponent of the incautious principle can restrict the scope of the principle so that a certain degree of confidence is necessary to trigger it. For example, they can stipulate that the incautious principle applies if and only if you are, say, less than 15% confident that a particular individual is sentient; otherwise one of the other principles that we will be considering here applies. This third reply strikes me as more plausible, and worth considering. Of course, we might wonder if this restriction has a principled basis. But we can ignore this complication for now.

However, even if we accept this third reply, the incautious principle still faces another, more important objection. Even if we restrict the scope of this principle to apply only below a particular degree of confidence in sentience, and even if we assume that our degrees of confidence are precise enough to make this kind of restricted principle work, it still has implausible implications regarding individuals who fall within its restricted scope. For example, suppose that you are only, say, 12% confident that a lobster is sentient (and therefore that this lobster is within the scope of the incautious principle). Suppose further that you feel inclined to boil this lobster alive – not even in order to eat them, but rather only for the simple pleasure of watching them boil alive. In this case, the incautious principle implies that you have *no moral duty to the lobster* not to act on this inclination, all else equal. However, I suspect that many of us disagree with this assessment. We think that, if there is a real possibility that this lobster is sentient, then you morally ought to take this possibility into account when thinking about how to treat them. And if you do not – for example, if

you boil them alive for the simple pleasure of doing so despite believing that there is a 12% chance that they are phenomenally consciously experiencing your boiling them alive – then your action is at least prima facie morally wrong. Moreover, what makes your action at least prima facie morally wrong is not only, as Kant claimed, that you might be conditioning yourself to harm other human and nonhuman animals, but also, indeed primarily, that you might be harming this particular lobster.¹¹

If this is right, then the lesson is: We need to find a solution to the sentience problem that makes our moral thinking sensitive to the possibility that other beings are sentient. The question is: How can we do that?

3.2. The precautionary principle

This brings us to *the precautionary principle*, which requires treating other beings as sentient in cases of uncertainty.¹² Many philosophers accept precautionary principles in other contexts, and some seem to accept a precautionary principle in this context as well.

What should we think of the precautionary principle in this context? We can start by noting that precautionary principle is more plausible than the incautious principle in the lobster case that we considered above, since it supports the idea that you have a moral duty to the lobster not to boil them alive for the simple pleasure of doing so, all else equal.

However, even if the precautionary principle is more plausible than the incautious principle in this kind of case, it still faces at least two objections, which resemble the two objections that the

¹¹ Immanuel Kant, *Lectures on Ethics*, ed. J. B. Schneewind. (Cambridge: Cambridge University Press, 2001).

¹² See Francione 2012 for a statement of this principle. Relatedly, Guerrero 2007 argues that you are blameworthy for killing individuals unnecessarily in cases of (subjective) uncertainty about whether or not they have moral status. I think that this principle is plausible, but I also think that is neutral between the precautionary principle, which I consider in this section, and the expected value principle, which I consider in the next section.

incautious principle faces. The first objection is that the precautionary principle, as currently stated, is too extreme. For example, suppose that you are open to the possibility that panpsychism is true, i.e. that everything in the world is phenomenally conscious, and you are also open to the possibility that functionalism about pleasure and pain is true, i.e. that the capacity to detect, and respond to, helpful and harmful stimuli is sufficient for the capacity to experience pleasure and pain. Finally, suppose that, as a result of these expressions of epistemic humility, you are also open to the possibility that a very wide range of beings in the world – insects, plants, robots, and more – are sentient (even if they might not be as sentient as, say, humans and other vertebrates are). In this case, the precautionary principle implies that you morally ought to treat a very wide range of beings in the world – insects, plants, robots, and more – as having moral status (even if you might not need to treat them as having the same kind or degree of moral status as, say, humans and other vertebrates). But one might object that this kind of moral animism is a non-starter as well. The mere fact that current generation smart phones are, say, .00001% likely to be sentient, given your evidence, is not sufficient reason to treat them as though they are.

As with the incautious principle, a proponent of the precautionary principle has three possible replies to this objection. First, they can deny that the precautionary principle leads to moral animism. For example, they can assert that they are, in fact, 100% certain that panpsychism is false, or that functionalism about pleasure and pain is false. In this case, the precautionary principle would not imply that they are required to treat, say, current generation smart phones as sentient. This reply strikes me as somewhat plausible, though I personally feel less than 100% certain about this.

Second, a proponent of the precautionary principle can deny that moral animism is a non-starter. For example, they can accept the idea that we have moral duties to a very wide range of beings in the world, and they can try to make this idea seem plausible by noting that, even if we have moral duties to a very wide range of beings in the world, we might still have stronger duties to, say,

human and nonhuman animals than they have to, say, current generation smart phones. This reply strikes me as somewhat plausible too, though perhaps too blunt in some cases in a way that I will explain below.

Third, a proponent of the precautionary principle can restrict the scope of the principle so that a certain degree of confidence is necessary to trigger it. For example, they can stipulate that the precautionary principle applies if and only if you are, say, more than 5% confident that a particular individual is sentient; otherwise one of the other principles that we are considering here applies. This third reply strikes me as somewhat plausible as well, though a lot depends on how we decide to treat individuals to whom the restricted precautionary principle does not apply. Additionally, as before, we might wonder if this restriction has a principled basis. But we can ignore this complication for now.

However, even if we accept this third reply, the precautionary principle still faces another, more important objection. Even if we restrict the scope of this principle to apply only above a particular degree of confidence in sentience, and even if we assume that our degrees of confidence are precise enough to make this kind of restricted principle work, it still has implausible implications regarding individuals who fall within its restricted scope. For example, suppose that you are, say, 12% confident that a lobster (hereafter bio-lobster) is sentient, and you are, say, 8% confident that a functionally identical robot lobster (hereafter robo-lobster) is sentient.¹³ Finally, suppose that a house containing both lobsters is burning down and you can save either but not both. Assuming you should save one of the lobsters, which one should you save? The precautionary principle implies that you are

¹³ Why these degrees of confidence? On one hand, the bio-lobster and robo-lobster are functionally identical, which is why your degrees of confidence are so similar. (This is not a current generation robot lobster, but rather a future generation robot lobster with artificial general lobster intelligence.) On the other hand, the bio-lobster is physiologically and evolutionarily continuous with you whereas the robo-lobster is not, and therefore you have at least a bit more reason to think that the bio-lobster is sentient than that the robo-lobster is. Of course, some readers might disagree with this analysis. If so, readers are welcome to replace these lobsters with other beings about whom they can imagine having roughly these degrees of confidence

morally required to treat both lobsters as sentient, and therefore it supports neutrality about which lobster you should save all else equal. Yet I suspect that many of us think that this assessment of the situation is too simple: We think that, while both of these lobsters *might* be sentient, the bio-lobster is also a bit more *likely* to be sentient than the robo-lobster is, given your evidence, and therefore you morally ought to break the tie in favor of saving the bio-lobster, all else equal.

3.3. Why the precautionary principle is more plausible than the incautious principle

As we have seen, the incautious principle and precautionary principle have symmetrical risks. The incautious principle has a higher risk of leading to false negatives (i.e., to mistakenly treating sentient beings as non-sentient), whereas the precautionary principle has a higher risk of leading to false positives (i.e., to mistakenly treating non-sentient beings as sentient). Moreover, proponents of these principles can respond to these risks in similar ways, that is, either by accepting this risk, by rejecting the risk, or by restricting their principle to reduce this risk. However, as I have indicated, these replies strike me as more plausible when made in support of the precautionary principle than made in support of the incautious principle. I will now say a bit more about why.

In general, when we need to decide between a principle that risks false positives and a principle that risks false negatives, we can decide which to accept by comparing the harm of false positives and the harm of false negatives in the relevant context. For example, many people believe that, in the context of criminal justice, false positives (i.e., mistakenly treating innocent people as guilty) are worse than false negatives (i.e., mistakenly treating guilty people as innocent) all else equal. As a result, we think that a criminal justice system should be designed so as to create a presumption of innocence rather than a presumption of guilt, all else equal.

In the context of the moral problem of other minds, the question is: How does the harm of mistakenly treating non-sentient beings as sentient compare with the harm of mistakenly treating sentient beings as non-sentient? In my view, false negatives are clearly worse than false positives in this context, since the harm that we cause by accidentally treating moral subjects as non-moral subjects is clearly worse than the harm that we cause by accidentally treating non-moral subjects as moral subjects. If this is right, then we morally ought to adopt a presumption of sentience rather than a presumption of non-sentience, all else equal.

Granted, one might argue that false positives are worse than they initially seem. For example, insofar as we mistakenly treat non-sentient beings as sentient, there is a risk that we will be wasting scarce moral resources on them. After all, humans and other vertebrates already need more help than they receive. Imagine how hard it will be for these individuals if we decide that quintillions of other individuals need help too. One might even worry that if we try to help everyone, then we will end up helping no one. For example, if you try to let everyone on a life boat, then the life boat will sink. Similarly, one might worry, if we try to distribute scarce moral resources across quintillions of possibly sentient beings, then nobody will have enough to survive or flourish.

This is a reasonable concern. However, we can respond to it in at least three ways. First, we can deny the premise. It might be that in the context of our current social, political, and economic systems, we would be spreading ourselves too thin if we tried to distribute moral resources across quintillions of possibly sentient beings. But maybe in the context of other social, political, and economic systems, we would not be spreading ourselves too thin if we tried to do that. Or at least, maybe we would not be spreading ourselves too thin if we tried to distribute moral resources to more beings than we currently do.

Second, we can accept the premise but deny the conclusion. Equal consideration does not always mean equal treatment. This is true in at least two respects. First, even if we extended moral

consideration to all possibly sentient beings, we might still think that there are morally relevant differences among them. For example, we might think that some individuals are more sentient than others, and therefore that we can be morally permitted or required to prioritize these individuals in some choice situations. We might also think that we have closer relationships with some individuals than with others, and therefore that we can be morally permitted or required to prioritize these individuals in some choice situations.

The second, related respect in which equal consideration does not mean equal treatment is that, even in cases where equal treatment would be best, we cannot always treat everyone equally. For example, we are not always (or maybe ever) able to help all the humans who need our help. Does that mean that we should deny that many humans merit moral consideration at all? For example, can we say, “Nobody else can fit on this life boat, so nobody else matters! We might as well throw things at people while they drown!” Of course not. Instead, it simply means that life is tragic, and we cannot always treat everyone as they deserve to be treated. Similarly, we might think, we are not always (or maybe ever) able to help the quintillions of nonhumans who need our help. Does this mean we should deny that many nonhumans merit moral consideration at all? Not necessarily. Instead, it might simply mean that life is even more tragic than we originally thought.

The third response that we can make to this concern about moral overextension is to look for a middle-ground principle that reduces risk of false negatives as well as of false positives. As we have seen, one way to do this is by restricting the scope of the incautious and precautionary principles. For example, we can use the precautionary principle if we are at least 5% confident that other beings are sentient, and we can use the incautious principle otherwise. Another way to do this is to accept a different kind of principle that allows for scalar judgments. For example, we can use a principle that allows us to rank other beings based on how likely they are to be sentient, given our evidence. This response would also have the bonus of implying that you should save the bio-lobster

over the robo-lobster in the burning house case (assuming, that is, that you think that the bio-lobster is more likely to be sentient than the robo-lobster, given your evidence).

In this case, the question that we face is: Can we find a way to make our solution to the moral problem of other minds sensitive to not only the possibility but also the (subjective) probability that other beings are sentient? And if so, would this scalar solution to the sentience problem be better or worse than a (perhaps suitably restricted) precautionary principle?

4. The expected value principle

This brings us to *the expected value principle*. This principle holds that, in cases of uncertainty about whether or not a particular being is sentient, we are morally required to multiply our degree of confidence that they are sentient by the amount of moral value that they would have if they were, and to treat the product of this equation as the amount of moral value that they actually have.

How does this kind of principle compare with the precautionary principle? It is hard to answer this question in the abstract, since the expected value principle will work differently in the context of different moral theories. In what follows I will consider utilitarian and Kantian interpretations, to illustrate what some of the main considerations will be. I will then argue that the precautionary principle and expected value principle (on utilitarian as well as Kantian interpretations) have familiar pros and cons, and that which we should use is likely to be a theory- and context-specific matter.

4.1. Utilitarianism

Utilitarians, who think that we morally ought to maximize pleasure and minimize pain in the world, will interpret the expected value principle as an *expected utility principle*. On this view, in cases of uncertainty about whether other beings are sentient, we are morally required to multiply our degree of confidence that they are by the amount of pleasure and pain they would be experiencing if they were, and then treat the product of this equation as the amount of pleasure and pain that they are actually experiencing.

We can start by noting that, if we can make this kind of principle work, then it will have different implications than the precautionary principle in many cases. For example, whereas the precautionary principle is neutral about whether to save the bio-lobster and the robo-lobster in the burning house case, the expected utility principle supports saving the bio-lobster all else equal. More generally, whereas the precautionary principle requires us to treat all possibly sentient beings as sentient, with the result that it will have more demanding implications regarding, say, insects and robots, the expected utility principle allows for a more scalar approach, with the result that it will have less demanding, though still quite demanding, implications regarding these individuals.

The expected utility principle also faces at least two objections, however. First, we might worry that we are not capable of making reliable comparative judgments about sentience in many cases. This is true whether we make these judgments reflectively or intuitively. For example, think about what it would take to make these judgments reflectively. We would have to calculate the probability that other beings are conscious, the probability that other beings experience pleasure and pain, the level of conscious pleasure and pain they would be able to experience if any at all, and more. And, we might worry that it would be difficult if not impossible for us to calculate these probabilities and utilities reliably in many cases.¹⁴

¹⁴ For discussion of how difficult intersubjective comparisons of utility are even on the assumption that the entities in question are sentient, see Gruen 2011.

Similarly, think about what it would take to make these predictions intuitively. We would have to make snap judgments about how likely each individual is to be sentient and then rank them accordingly. But as before, we have good reason to doubt that these snap judgments are reliable in many cases. For example, studies show that our capacity for empathy is sensitive to a variety of factors that we think, on reflection, are not likely to be relevant, including big vs. small eyes, furry vs. slimy skin, four vs. six limbs, symmetrical vs. asymmetrical features, and so on.¹⁵ Thus, we might worry that, if we try to make comparative judgments about sentience in many cases, then we may well end up doing more harm than good as a result.¹⁶

A utilitarian has at least two possible replies to this objection. First, they can argue that we can, in fact, make reliable comparative judgments about sentience in many cases. For example, in the burning house case, you do not need to assign an exact probability to the proposition that each lobster is sentient. Instead, all you need to do is break a tie between them. And, given that these lobsters are identical in every respect, except that the bio-lobster is physiologically and evolutionarily continuous with you whereas the robo-lobster is not, you might think that a decision to break the tie in favor of the bio-lobster is, at the very least, likely to be better than chance.¹⁷

Second, a utilitarian can accept that we cannot make reliable comparative judgments about sentience in many cases, and that we should not use an expected utility principle in these cases as a result. For example, suppose that, instead of a bio-lobster and a robo-lobster, the burning house contains a robo-lobster and a bio-Roomba (where a bio-Roomba is a simple organism who wakes up,

¹⁵ For discussion of the many ways in which snap judgments about sentience can mislead, see Hal Herzog 2010. For a critique of anthropomorphism in cognitive ethology and comparative psychology, see Wynne 2007. For a defense of anthropomorphism in these fields, see Fisher 1996.

¹⁶ For an argument that expected utility analysis concerning the distant future might do more harm than good overall, see Jamieson 1992.

¹⁷ For an empirically and philosophically rigorous (but nevertheless tentative) attempt to assign probabilities of sentience to different beings, see Muehlhauser 2017.

grazes on dust, and then falls asleep again).¹⁸ Assuming you should save one of these individuals, which one should you save? In this case, you might not have enough information to be able to make a reliable comparative judgment about sentience (especially given how many differences there are between these individuals and how little time you have to make a decision). If so, a utilitarian might grant that you should not use an expected utility principle to decide what to do. Instead, they might say that you should use a precautionary principle and then find a different basis for making a decision (or perhaps simply decide randomly).

Even if we accept these replies to the first objection, however, we still face a second objection, which is that the expected utility principle is unacceptably biased. We might think that we should have a higher degree of confidence that beings who are behaviorally, physiologically, and evolutionarily continuous with us are sentient than beings who are not, all else equal. If so, then the expected utility principle will imply that we should morally prioritize the former beings over the latter, all else equal. For example, in a future world where we all co-exist with functionally identical robots, the expected utility principle might direct us to systematically favor bio-humans to robo-humans all else equal (and vice versa). And we might think that this kind of discrimination is bad for epistemic reasons (it might be mistaken) as well as for practical reasons (it might be socially harmful).

In response to the epistemic concern – that this kind of discrimination might be mistaken – a utilitarian has three possible responses. First, they can deny that the expected utility principle has this implication. For example, if we are 100% certain that functionalism is true, then perhaps we can assign the same probability of sentience to bio-humans and robo-humans. As above, I think that this reply is somewhat plausible, though I have my doubts. Second, a utilitarian can accept that the expected utility principle has this implication, and they can reject the expected utility principle in

¹⁸ Thanks to B.J. Rosen for suggesting this case.

some cases as a result. That is, if we are sufficiently concerned about the risk of bias in some cases, then we can correct for that risk by not making scalar judgments at all, or at least by not making them based on relations of similarity and difference. Third, a utilitarian can accept that the expected utility principle has this implication, and they can accept this implication in some cases as a result. Granted, what maximizes expected utility is not always the same as what maximizes actual utility. But at the end of the day, we have no choice but to do the best we can with the evidence we have available.

In response to the practical concern – that this kind of discrimination might be socially harmful – a utilitarian once again has three possible responses. First, they can deny that the expected utility principle has this implication, for the same kind of reason as above. Second, they can accept that the expected utility principle has this implication, and they can reject the expected utility principle in some cases as a result. That is, if we are sufficiently concerned about the risk of social conflict in some cases, then we can correct for this risk by not making scalar judgments at all, or at least not making them based on relations of similarity and difference.¹⁹ Third, a utilitarian can accept that the expected utility principle has this implication, and they can accept this implication in some cases as a result. Granted, insofar as the expected utility principle risks causing social conflict, this would be bad. But insofar as it at least does more good than harm relative to alternatives, most utilitarians will still be prepared to accept it all things considered.

4.2. Kantianism

In contrast, Kantians, who think that we morally ought to treat moral subjects as ends in themselves, will interpret the expected value principle as an *expected dignity principle*. On this view, in cases of uncertainty about whether other beings are sentient, we are morally required to multiply our degree

¹⁹ For discussion of indirect morality, see Hare 1982.

of confidence that they are by the absolute and incomparable moral value that they would have if they were, and to treat the product of this equation as amount of moral value that they actually have.

As with the expected utility principle, we can start by noting that, if we can make this kind of principle work, then it will have different implications than the precautionary principle in many cases. For example, it will support saving the bio-lobster over the robo-lobster in the burning house case²⁰, and it will also allow for a scalar approach to priority-setting more generally, with some exceptions. The question here is: Can we make this kind of principle work?²¹

We might worry that the answer to this question is no, since, we might think, the expected dignity principle is not meaningfully different from a Kantian precautionary principle. Why not? Consider how Kantians think about moral status. If moral status is absolute and incomparable, then it must also be infinite. Yet if moral status is infinite, then the expected dignity principle implies that we should treat all possibly sentient beings as having full and equal moral status – since, if we multiply any non-zero number by infinity, the product is infinity. Thus, for example, we might think that the expected dignity principle implies that you should treat both lobsters as having full and equal moral status in the burning house case, and therefore, like the precautionary principle, it supports neutrality about which lobster you should save all else equal.

A Kantian might reply to this worry, however, by denying that moral status must be infinite in order to be absolute and incomparable. Yes, we can think of moral status as infinite if doing so helps us to treat moral subjects as ends in themselves (rather than as mere contributors to aggregate

²⁰ Many Kantians view helping as an imperfect duty that applies in some but not all cases. We can accommodate this in the burning house case by imagining that you are responsible for the fire, and therefore that your decision about which lobster to save is not a decision about which lobster to *help*, but rather a decision about which lobster to *avoid harming*.

²¹ Some people might wonder why we are considering Kantianism at all here, since Kant believed that you have moral status if and only if you are rational. However, philosophers such as Korsgaard 2011 and Regan 2004 have challenged this interpretation of Kantianism, and have developed sentientist Kantianism as an alternative to rationalist Kantianism. My discussion in what follows will engage with this contemporary, sentientist interpretation of Kantianism.

utility). But there is no need for us to take this idea of infinity literally. As long as we treat moral subjects as ends in themselves, we can attach any arbitrary finite value to moral status for purposes of doing risk analysis in certain cases. For example, we can stipulate that full moral status is equivalent to 100 units of value, not because we think that moral subjects actually have 100 units of value (whatever that would even mean for a Kantian), but rather simply because we need a finite number to use in cases where our degrees of confidence about sentience are morally relevant.

For example, imagine that two houses are burning down, and you can save either but not both by pulling a lever that will direct water to that house. The only information you have about these houses is: The house on the left has lights on, whereas the house on the right has lights on *and* music playing. Assuming you should save one of these houses, which one should you save?²² It seems consistent with Kantianism to say that you morally ought to save the house on the right all else equal. Why? Because even though anyone who might be home has full and equal moral status, someone is a bit more *likely* to be home on the right, given your evidence, and therefore you minimize risk of harm to individuals with full and equal moral status if you save the house on the right.

If this is right, then a Kantian can interpret the expected dignity principle in the same kind of way. On this interpretation, the question is not whether we should treat some beings as having a higher moral status than others, but rather whether we should treat some beings as having a higher chance of “having someone home” – i.e. of having sentience that confers full and equal moral status – than others in certain cases. For example, in the original burning house case, a Kantian can say that you morally ought to save the bio-lobster all else equal. Why? Because even though anyone who might be “home” in these lobsters has full and equal moral status, someone is a bit more *likely* to be

²² As above, many Kantians view helping as an imperfect duty that applies in some but not all cases. We can accommodate this by imagining that you are responsible for the fire, and therefore that your decision about which house to save is not a decision about which people to *help*, but rather a decision about which people to *avoid harming*.

home in the bio-lobster, given your evidence, and therefore you minimize risk of harm to individuals with full and equal moral status if you save the bio-lobster.

If this is right, then the expected dignity principle can be a meaningful alternative to a Kantian precautionary principle after all. At this point, however, one might worry that the expected dignity principle inherits not only the pros but also the cons of the expected utility principle (with some differences, which I will discuss in a moment). Yes, the expected dignity principle allows us to take a scalar approach in some cases of uncertainty about sentience. At the same time, it raises questions about cluelessness, since it requires us to make comparative judgments that we might not always be able to make reliably, as well as questions about bias, since it requires us to favor individuals who are like us in certain respects, which might be bad for epistemic as well as practical reasons.

Can Kantians reply to these worries in the same way that utilitarians can? For the most part, I think that they can. Like utilitarians, Kantians can use the expected dignity principle in cases where doing so seems epistemically and practically best (as well as otherwise morally permissible), and they can use the precautionary principle otherwise. However, there will be differences. For example, utilitarians tend to have a higher tolerance for cluelessness, demandingness, and other such challenges than Kantians. So, insofar as the expected value principle raises these kinds of concerns, utilitarians might be more willing to accept this principle anyway. Similarly, utilitarians and Kantians will have different standards for when using this principle is morally permissible. I will now say a bit more about these differences.

4.3. Why the precautionary and expected value principles have familiar pros and cons

If my discussion so far is correct, then it leaves us in a familiar place. The precautionary principle

and expected value principle have the same kinds of pros and cons in this context that they have in other contexts, for utilitarians as well as Kantians (and other moral theorists). The precautionary principle raises more concerns about demandingness whereas the expected value principle raises more concerns about cluelessness. Meanwhile, the expected value principle works similarly for utilitarians and Kantians in some respects (e.g., utilitarians and Kantians can rank individuals by probability of sentience for some purposes, such as deciding whom to save) but differently in other respects (e.g., utilitarians but not Kantians can rank individuals by probability of sentience for other purposes, such as deciding how to aggregate wellbeing across individuals). The upshot is that which principle we should use is likely a theory- and context-specific matter.

Consider each of these points in turn. First, the precautionary principle and expected value principle both raise concerns about demandingness, but the precautionary principle raises these concerns a bit more. Why? They both raise concerns about demandingness because they both require us to treat possibly sentient beings as at least partially sentient. However, the precautionary principle raises these concerns a bit more than the expected value principle because it requires us to treat possibly sentient beings as sentient, whereas the expected value principle permits us to treat them as partially sentient, depending on the probabilities. For example, suppose that we think that there is a 5% chance that insects are sentient, and that the average insect, if sentient, experiences 0.1% the amount of conscious pleasure and pain as the average human at any given time. On a precautionary principle with no (or sufficiently low) scope restriction, we would have to say that the total amount of conscious pleasure and pain experienced by 10 quintillion insects is equal to that of 10 quadrillion humans. In contrast, on an expected value principle with no (or sufficiently low) scope restriction, we would be able to say that the total amount of conscious pleasure and pain experienced by these insects is equal to that of only (!) 500 trillion humans. In this case, then, both principles raise

concerns about how demanding our moral duties to insects might be, but the precautionary principle does so more than the expected value principle.

With that said, a lot depends on which moral theory we accept. For example, utilitarians think that we should maximize aggregate wellbeing in the world. Thus, the precautionary principle and expected value principle (especially the former) will have more demanding implications for utilitarians on a global scale. For example, on a utilitarian view, if the total amount of invertebrate pleasure and pain in the world is greater than the total amount of vertebrate pleasure and pain in the world, then we should prioritize invertebrates over vertebrates all else equal. In contrast, since Kantians deny that we should maximize aggregate wellbeing in the world, these principles may not have this kind of implication for them, depending on the details (though, if we are complicit in harm to invertebrates on a global scale, then that might change things for Kantians). Meanwhile, many Kantians think that we should avoid harming or killing moral subjects merely as means to our ends, or as an unnecessary side effect of pursuing our ends. Thus, the precautionary principle and expected value principle (especially the former) will have demanding implications for Kantians on a local scale. For example, on a Kantian view, insofar as normal activities such as lawn care involve harming or killing insects as means to our ends, or as an unnecessary side effect of pursuing our ends, these activities will be morally wrong all else equal. In contrast, since utilitarians deny such restrictions on harming or killing, these principles may not have this kind of implication for them, depending on the details (though, if the activities that cause local insect suffering are not necessary for maximizing utility, then that might change things for utilitarians).

Second, the precautionary principle and expected value principle both raise concerns about cluelessness, though the expected value principle raises these concerns a bit more. Why? They both raise concerns about cluelessness because they both require us to estimate whether or not other beings are possibly sentient. They also, if we accept a scope restriction for the precautionary

principle, require us to estimate probability of sentience at least in some cases (e.g., when determining whether or not a particular being is within the scope of a restricted precautionary principle). However, the expected value principle raises these concerns more than the precautionary principle because it requires us to estimate probability of sentience in a wider range of cases than the precautionary principle does. All the precautionary principle requires us to do is decide whether or not the probability that a being is sentient is above a particular baseline. In contrast, the expected value principle requires us to make comparative judgments about sentience regarding a wide range of beings in a wide range of contexts. As we have seen, this is hard to do well for both scientific and philosophical reasons. Thus, both of these principles raise concerns about whether we have enough information to reliably follow them, but the expected value principle does so more than the precautionary principle.

With that said, a lot depends on which moral theory we accept here too. For example, since utilitarians think that we should maximize wellbeing in the world, they will think that we should try to make comparative judgments about sentience in many cases. Can we do more good by helping or not harming five humans or 50 dogs? 50 dogs or 5,000 lobsters? 5,000 lobsters or 5,000,000 ants? And so on. Insofar as we can calculate these probabilities and utilities reliably (and without causing other kinds of harm), a utilitarian will say that we should do so all else equal. Otherwise we should not. In contrast, since Kantians deny that we should maximize wellbeing in the world, they will not think that we should try to make comparative judgments about sentience in quite as many cases. Yes, we might have reason to make a comparative judgment about sentience in order to decide whether to save the bio-lobster or the robo-lobster (since this decision does not involve aggregation, and does not involve harming anyone as a means to our ends, or as an unnecessary side effect of pursuing our ends). But we do not have reason to make comparative judgments about sentience in order to decide whether or not to help or not harm five humans, 50 dogs, 5,000 lobsters, and so on (since this would

involve aggregation and, perhaps, impermissible harm). If this is right, then the kinds of cases in which the expected utility principle is an option for utilitarians plausibly raise cluelessness concerns than the kinds of cases in which this principle is an option for Kantians.

Finally, the precautionary principle and expected value principle have competing pros and cons with respect to general intuitive plausibility. On one hand, as we have seen, the precautionary principle might seem to be too non-scalar in some cases. If we have to choose between helping or harming two beings, and the only difference is that we have more reason to believe that one of these beings is sentient than that the other of these beings is sentient, then it seems plausible that we should break the tie in favor of the former being rather than the latter being. The expected value principle accommodates this idea, whereas the precautionary principle does not. On the other hand, as we have seen, the expected value principle might seem to be too scalar in some cases. We should be wary about the epistemic and practical risks involved with creating hierarchies based on relations of similarity and difference. The precautionary principle accommodates this idea, whereas the expected principle does not. (Though of course, even if we accept the precautionary principle, we still face questions about whether to rank beings on other grounds, such as the kind or degree of sentience that we take them to have, if any at all.)

A lot depends on which moral theory we accept here too, for the same reasons as above. For example, utilitarians have a relatively high tolerance for demandingness, cluelessness, and intuitive implausibility. They are also willing to accept *indirect* and *esoteric morality* if need be, i.e. to accept that, if we can do more good by following or promoting incorrect moral principles than we can by following or promoting correct ones, then we morally ought to do so. In contrast, Kantians (and other moral theorists) have a relatively low tolerance for demandingness, cluelessness, implausibility, and indirect or esoteric morality. As a result, insofar as the precautionary and expected value principles have the tradeoffs we have considered here, Kantians (and other moral theorists) have a harder

choice to make. In particular, they will need to either (a) find a principled basis for restricting the scope of these principles, (b) find a principled basis for harming or not helping quintillions of moral subjects, or (c) accept that morality involves more demandingness, cluelessness, or intuitive implausibility than they originally thought. My own view is that Kantians (and other moral theorists) should accept a combination of (b) and (c). I will not argue for this view here, except to say that, insofar as one accepts (a), one will have to explain how this view is different from a restricted incautious principle – or why a restricted incautious principle is rationally or morally acceptable.

5. Conclusion

I think that my discussion in this paper supports three conclusions about the moral problem of other minds.

First, the precautionary and expected value principles are more plausible than the incautious principle. If there is a chance that a particular individual is sentient, then we morally ought to take that into account when thinking about how to treat them. It follows that, unless and until other, better solutions to the moral problem of other minds become available, we should accept either the precautionary principle, the expected value principle, or both. For example, we can use the expected value principle in cases where we have enough information to do so reliably (and where doing so is otherwise morally permissible), and we can use the precautionary principle in cases where we do not have enough information to use the expected value principle reliably (or where doing so is otherwise morally impermissible).

Second, if we accept the precautionary principle, the expected value principle, or both, then we morally ought to treat many beings, including but not limited to nonhuman animals, as having at least partial moral status. Depending on the details, this list might include fetuses, animals, plants,

robots, and more. Of course, as we have seen, to say that we should treat these beings as having at least partial moral status is not to say how we should treat them all things considered. For example, Judith Jarvis Thomson argues that abortion can be morally permissible in many cases even if fetuses have moral status.²³ Similar analyses will apply to the ethics of harming or killing insects, robots, and so on. Still, if this conclusion holds, it will raise moral questions about many common practices.

Third, if we morally ought to treat many beings, including but not limited to nonhuman animals, as having at least partial moral status, then morality involves more cluelessness and demandingness than we might have thought. As we have seen, however, a lot will depend on the details. For instance, since utilitarians accept aggregation, they will have to accept more revisionary conclusions about global issues, such as the relative moral priority of vertebrates and invertebrates all else equal. In contrast, since Kantians reject aggregation, they will have to accept more revisionary conclusions about local issues, such as the permissibility of harming insects as a means to lawn care. Either way, then, we will likely find that a commitment to the precautionary and/or expected value principles will shape and constrain our choices in life in surprising and profound ways.²⁴

6. References

Block, Ned.. "On a confusion about the function of consciousness." *Behavioral and Brain Sciences* 18 (1995) : 227–47.

Braithwaite, Victoria. *Do Fish Feel Pain?* Oxford: Oxford University Press, 2010.

Carruthers, Peter. "The Problem of Other Minds." In *The Nature of the Mind: An Introduction*, 6-35. New York: Routledge, 2004.

²³ Judith Jarvis Thompson, "A Defense of Abortion," *Philosophy and Public Affairs* 1, no. 1 (1971).

²⁴ Thanks to the participants at the 8th Rocky Mountain Ethics Congress, the Bioethics Works in Progress Series at the National Institutes of Health, the PPE and Value Theory Workshop at the University of North Carolina at Chapel Hill, and the Philosophy Colloquium at the University of South Carolina for helpful feedback and discussion. Thanks also to the editors of this journal for useful feedback and to Dale Jamieson for detailed discussion throughout the writing process.

DeGrazia, David. *Taking Animals Seriously: Mental Life and Moral Status*. Cambridge: Cambridge University Press, 1996.

Fisher, John. "The Myth of Anthropomorphism." In *Readings in Animal Cognition*, ed. Colin Allen & D. Jamieson, 3-16. Cambridge: MIT Press, 1996.

Francione, Gary. "Sentience," *Animal Rights: The Abolitionist Approach*, posted July 12th, 2012. Accessed December 15 2016. <http://www.abolitionistapproach.com/sentience/>

Gruen, Lori. "Experimenting with Animals." In *Ethics and Animals*, 105-129. Cambridge: Cambridge University Press, 2011.

Guerrero, Alex. "Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution," *Philosophical Studies* 136 (2007): 59–97.

Hansson, Sven Ove. "Risk," *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, 2014. Accessed December 15, 2016, <http://plato.stanford.edu/archives/spr2014/entries/risk/>.

Hare, Richard. *Moral Thinking: Its Levels, Methods, and Point*. Oxford: Oxford University Press, 1982.

Herzog, Hal. "The Importance of Being Cute." In *Some We Love, Some We Hate, Some We Eat*, 37-66. New York: Harper Collins, 2010.

Huebner, Bryce. *Macrocognition: A Theory of Distributed Minds and Collective Intentionality* Oxford: Oxford University Press, 2004.

Jamieson, Dale. "Ethics, Public Policy, and Global Warming," *Science, Technology, and Human Values* 17, no. 2 (1992): 139-53.

Kant, Immanuel. *Lectures on Ethics*, ed. J. B. Schneewind. Cambridge: Cambridge University Press, 2001.

----- *The Metaphysics of Morals*, ed. Mary Gregor. Cambridge: Cambridge University Press, 1996.

Korsgaard, Christine. "Interacting with Animals: A Kantian Account." In *The Oxford Handbook of Animal Ethics*, ed. Tom L. Beauchamp and R. G. Frey, 91-118. Oxford: Oxford University Press, 2011.

Muehlhauser, Luke. "2017 Report on Consciousness and Moral Patienthood," Open Philanthropy Project, 2017. Accessed March 20, 2018: <https://www.openphilanthropy.org/2017-report-consciousness-and-moral-patienthood>

Nozick, Robert. *Anarchy, State, and Utopia*. New York: Basic Books, 1984.

Regan, Tom. *The Case for Animal Rights*. Berkeley: The University of California Press, 2004.

Scheffler, Samuel. *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. Oxford: Oxford University Press, 1984.

Singer, Peter. *Animal Liberation*. New York: HarperCollins, 2009.

Thomson, Judith Jarvis. "A Defense of Abortion," *Philosophy and Public Affairs* 1 no. 1 (1971): 47–66.

Wallace, David Foster. "Consider the Lobster," *Gourmet Magazine* (August 2004).

Wynne, Clive. "What are Animals? Why Anthropomorphism Is Still Not a Scientific Approach to Behavior," *Comparative Cognition and Behavior Reviews* 2 (2007): 125-135.