



The Role of Neuroscience in Precise, Precautionary, and Probabilistic Accounts of Sentience

13

Adam J. Shriver

Abstract

Given that there is currently no consensus as to exactly which animals are sentient, how should we make moral decisions when we are uncertain as to which of the animals influenced by that decision are sentient? And how relevant is evidence from the neurosciences for making these decisions? In this chapter, I outline three different approaches toward incorporating uncertainty about sentience into moral decision-making: what I call precise, precautionary, and probabilistic approaches to sentience. I suggest that neuroscientific evidence has different relevance for each of these accounts. Precautionary approaches should be adopted to provide basic protections for animals even when we are uncertain about their sentience, but probabilistic accounts are more relevant for decisions where we need to carefully weigh positive and negative consequences of different possible decisions. Precise accounts can be useful for providing guidance but are not directly relevant for making decisions or guiding policy.

Keywords

Sentience · Consciousness · Arguments by analogy · Precautionary principle

The Fire: You have a choice. A fire has ignited at one of the local university's agricultural research facilities. The humans have all been evacuated but some of the animals remain behind for what is, without intervention, certain death. You are leading a team of volunteer firefighters driving a truck to the affected area and only have limited time. You can either save a barn full of ten cows, another barn full of fifteen chickens, two aquariums full of fifty fish, or a terrarium that contains seventy ants. All of the animals would, if saved, live exactly two more months of a relatively comfortable, though mostly bland, existence. Those who are left behind will die via asphyxiation (even the fish). As luck would have it, one of your colleagues

A. J. Shriver (✉)
University of Oxford, Oxford, UK
e-mail: adam.shriver@philosophy.ox.ac.uk

© The Author(s) 2020
L. S. M. Johnson et al. (eds.), *Neuroethics and Nonhuman Animals*, Advances in Neuroethics, https://doi.org/10.1007/978-3-030-31011-0_13

221

brought along the most recent edition of *Principles of Neural Science* and suggested that you quickly scan the pages to see what can be found out about the comparative neuroanatomy of the different species. What should you do? And should the info from the book, assuming that you'll have time to read it on the way, be relevant to how you make your decision?

This scenario offers a stark choice that seems at first glance to depend on how we value the lives of different types of animals. Though it clearly is not a very realistic scenario, it resembles other types of value judgments made daily that have consequences in the real world. For example, in invasive animal research, researchers often interpret the “replace” principle in the 3Rs guidelines (replace, reduce, refine) as suggesting that it is better to run an experiment on larger numbers of mice than it is to run a similar experiment on equal or even smaller numbers of monkeys. In food ethics, some have argued that we should shift to a diet of insects to avoid the welfare problems associated with modern industrial farming of birds and mammals [1]. Others have argued that animal activists should focus on improvements to fish welfare compared with that of large mammals and even chickens due to the huge numbers of fish killed every year [2]. Each of these recommendations would seem to depend upon some method for comparing the lives and welfare of different species.

But though it might seem as though these different perspectives all require different evaluations of what types of capacities ground moral standing in animals, this need not be the case. All of the above positions can be reached from a starting point that assumes that all conscious experiences of positive and negative feelings are morally significant. Though two welfare promoters may agree that sentience matters for moral standing, they may nevertheless disagree about how to weigh potential evidence for sentience. Specifically, they may attach different significance to the extent to which different species have neural similarities with humans or other species they take to be clearly conscious, which in turn may determine at least in part how the potential interests of different species are weighed against one another.

In what follows, I examine the question of how we should treat neurological similarity under three different ways of thinking about sentience: what I call precise, precautionary, and probabilistic accounts of sentience. I argue that evidence from the neurosciences has very different importance for these three different types of accounts, to the point that evidence may not be required at all with regards to certain policy decisions. I suggest that each of the different accounts of sentience is appropriate for some contexts but not others. But prior to getting to these three accounts, I first need to say more about why assessing sentience and other mental capacities is such a challenge.

13.1 The Problem of Other Minds

In his *Meditations on First Philosophy*, Rene Descartes employed what some have called the Method of Doubt [3] in an attempt to find indubitable truths. He found that, when putting himself in a skeptical frame of mind, he could question almost everything he had previously taken for granted, including his knowledge of the existence an external world. However, due to his direct acquaintance with his own experience, he

believed there was one thing he could not seriously doubt, which led to the famous *cogito*: “I think, therefore I am.” Some have later questioned whether one could truly infer the existence of a continuing “self” that perseveres throughout our experiences [4], but nevertheless Descartes’ observation that there is something special about our direct acquaintance with our own mental experiences, and his attempts to build upon that observation, has remained a powerful example of philosophical thought.

Unfortunately, however, this methodology does not reveal much about the minds of others (or, arguably, even our own). We do not have direct, immediate access to anyone else’s mind. When trying to assess whether other humans are also conscious, the best we can do is to make inferences based either on how similar they are to us or other observations we suspect would be associated with conscious experience. This leads to the so-called Problem of Other Minds [5] or the problem of knowing that other humans are conscious.

This “problem” has been discussed extensively in the philosophical literature, but even philosophers who discuss it in great detail do not act in everyday life as if other people are insentient. And the very discussion can be enough to make some people roll their eyes at a certain style of philosophical methodology . . . surely no one could really take such concerns seriously! Similarly, many think that we should have an equally dismissive attitude to this type of skepticism when it is aimed at sentience in certain nonhuman animals. Can we really seriously doubt that companion cats and dogs are sentient even as they routinely exhibit clever ways of communicating their beliefs and desires? And if we think this about cats and dogs, it is hard to coherently think otherwise about many other kinds of animals with very similar behavior and neuroanatomy.

The problem, however, is that for some animals, the differences both in behavior and brain function are too dramatic to ignore. What should we think about animals that do not have anything resembling a brain or central nervous system and that do not respond to noxious stimulation or threats in the way we would expect? Many people draw the line once a certain limit is reached with respect to differences from humans, whether the line is drawn at fish, insects, or some other taxa or species. However, even when a line is drawn, given the extensive variation observed among biological organisms, can we come up with a principled explanation for justifying the differential treatment between those species on either side of the line?

Some have tried to sidestep this challenge by endorsing panpsychism, the view that phenomenal consciousness is a property that exists in all physical matter. However, the move from this position to specific moral recommendations is often unjustified. Even if we think that consciousness is ubiquitous, we know from our own experiences that not all conscious experience is morally significant. Some experiences are neutral and do not seem to have any positive or negative valence [6]. In fact, in rare circumstances, there appear to even be cases of pain that lack the unpleasantness that typically accompanies them [7]. So even if we were told definitively that, say, a hydrogen molecule has the property of consciousness to some degree, we would have no way of assessing whether that conscious experience was positive, negative, or entirely neutral. So, endorsing panpsychism just changes the challenging question from “how do we know which entities are conscious?” to “how

do we know which entities have morally significant consciousness?" Endorsing panpsychism takes us no closer to knowing which animals should be taken into consideration in our moral judgments.

So, we are left to confront a situation where, although there are various theories of consciousness, any honest assessment would have to conclude that we do not yet know what physical processes are most central to the experience of morally relevant feelings (or other psychological states). And it is not just that we do not know what physical processes are relevant; we do not even know what exactly it would look like to have sufficient evidence to know that certain physical processes are central.

But we do not have the luxury of throwing up our hands and embracing agnosticism. Humans have impacted and continue to impact potentially sentient organisms at a scale that is almost impossible to exaggerate. In fact, we have the capacity to end all sentient life on earth, but sadly the jury is still out as to whether we have the ability to organize and guide societies in a manner that would ensure that this does not happen. Meanwhile, everyday activities such as modern industrial agriculture are likely causing suffering to tens of billions of animals every year [8]. The stakes are high. We need some principled methodology for assessing sentience that can be of use in attempting to formulate ethical and legal restrictions on our treatment of nonhuman animals. I now turn to several different strategies that have been used to do this, and in particular will examine the relevance of neuroscience for making final decisions on these frameworks.

13.2 Precise Accounts of Sentience

Until recently, most philosophical attempts at using sentience to guide policy attempted to set fairly clear lines that could be of use for ethical guidelines. That is, criteria were provided that could be used to say, for any given species, whether that species should be treated as sentient or not. There was, and currently is, no consensus on the exact criteria, so the designation "precise" should not be thought of as indicating the *success* of the endeavors so much as the fact that they, in contrast to the other two groups of theories I will be considering, *aspire* to provide precise criteria that can be used to make the relevant determinations. For these theories, we can say that, *if* they are true, then they set decision criteria that can be used to determine which beings are sentient.

Gary Varner [9, 10] has suggested that most standard arguments about sentience in animals can be usefully interpreted as arguments by analogy. An exchange between Varner and Colin Allen about such arguments nevertheless nicely illustrates some general challenges facing precise accounts of sentience.

To begin, we can follow Varner in suggesting the following general format for an argument by analogy for animal sentience:

- P1. Things like X are known to have qualities a, b, c, . . . and m as well as n.
- P2. Things like Y are known to have qualities a, b, c, . . . and m.
- C. So, things like Y probably also have quality n. [9, p. 108]

Though most authors Varner is discussing do not explicitly formulate their arguments in these terms, they nevertheless can easily be represented in this way. Varner himself initially summarized four different meta-reviews of animal sentience to come up with a list of evidentiary conditions that may indicate sentience. This list includes both neural and nonneural criteria. The former include the presence of nociceptors, the presence of brains, the fact that nociceptors are connected to brains, and the fact that the organism has endogenous opioids. Among the latter criteria are behavioral responses to pain that resemble those of humans and behavioral responses to known analgesics. Taxa that ticked all of the boxes were regarded as good candidates for sentience. Based on these criteria, Varner initially concluded that most vertebrates were likely to be sentient, and invertebrates (with the exception of cephalopods) likely insentient.

However, as Allen [11] has pointed out, there are two serious weaknesses with using these arguments to make determinations about sentience. First, such assessments have to be based on the most current research, and the general trend has moved in the direction of discovering that in certain species capacities are present that were previously assumed not to exist. For example, in the 2 years between the time of Varner's initial argument and Allen's critique, nociceptors were discovered in several new species.

This concern can perhaps be addressed with what Varner describes as "going to war with the army you've got" [10, p. 115]. Surely all can agree that in the absence of conclusive evidence, we nevertheless have to make decisions based on what evidence we do have. However, as we will see, the type of line-drawing recommended by Varner can be problematic in ways that do not apply to other approaches I will consider.

Allen's second criticism reveals a deeper challenge. As he notes, just as there will be some biological similarities between humans and other animals, there will also always be some biological differences and in particular some neurological differences related to the capacity to feel pain and/or pleasure and/or other valenced states. So, sentience skeptics can point to these dissimilarities to argue against the claim that any given nonhuman animal is sentient. How, then, can we sort through the various similarities and dissimilarities to come up with the right criteria for comparison? Allen argues that we need a *guiding theory* that explains *why* we should expect certain features to be required for conscious experience that can guide the comparisons to firmer footing.

It is hard to argue against the suggestion that having such a theory would be ideal for any proposed account, and many authors have attempted to provide one, including Varner [10] in his response to Allen. Guiding theories for precise accounts of sentience might rely heavily on neuroscientific details or may eschew them entirely. But all guiding theories, as suggested by Allen, need to tell some story about why particular features give rise to conscious experience.

The problem, of course, is that there is intense disagreement about which theories should be preferred, and there is nothing coming close to compelling evidence to suggest that any particular theory is the right one. Meeting Allen's challenge with not just *a* guiding theory but rather *the correct* guiding theory requires nothing less

than discovering the neural correlates of sentience. In the absence of such a theory, we could be skeptical that humans are sentient, too. But, as stated above, there are many ethical and political decisions to be made, and waiting (perhaps indefinitely) for the correct guiding theory to make these decisions is likely to result in the suffering of numerous sentient beings.

So, it seems that we are left with Varner's "going to war with the army we have" concession. But if we adopt this approach, despite tenuous knowledge and disagreement about what the relevant traits truly are, what happens if we draw the line at the wrong place? If a precise account of sentience tells us we should draw the line at vertebrates, and our regulations are correspondingly formulated such that only vertebrates are protected (as is currently the case in the USA), what are the ethically significant consequences if we are mistaken? Consider, for example, invertebrate species such as lobsters that are routinely boiled alive for human food. If they are in fact capable of feeling pain, then treating them otherwise would be causing vast amounts of morally relevant suffering. As such, using a firm line to determine our ethical obligations seems likely to result in seriously bad consequences. But other approaches for dealing with evidence for sentience have been developed that can avoid this problem.

13.3 Precautionary Accounts of Sentience

Whereas precise accounts of sentience can be thought of as drawing lines based on what could be characterized as the *best* account of the relevant evidence for sentience, a precautionary approach to sentience might be loosely described as drawing lines, in certain conditions, based on a *sufficiently reasonable* account of relevant evidence for sentience. The precautionary principle, which in its original form stated "where there are threats of serious or irreversible damage, lack of full scientific certainty shall not be used as a reason for postponing cost-effective measures to prevent environmental degradation" [12, p. 3], was developed specifically to prevent serious harms from occurring in circumstances where scientific certainty is lacking. In particular, as Steele [13] notes, such a principle is needed to ensure that policy is not "susceptible to paralysis by scientific uncertainty" (p. 213). The principle is relevant for many challenges at the intersection of science and public policy, and perhaps most urgently with regard to climate change, where a number of interesting scientific questions remain despite it having been known for many years that action is warranted and, in fact, urgently needed.

The precautionary principle has been applied to nonhuman animals in the past, and a recent paper by Birch [14] provides a valuably detailed framework for applying it to policies that affect nonhuman animals. Birch begins by adapting the original policy-related precautionary principle so that it can be formulated to apply to nonhuman animals. The original precautionary principle references environmental degradation as the relevant harm, but if we think that other types of harms can be treated similarly, we can accept Birch's formulation of the Animal Sentience Precautionary Principle as: "Where there are threats of serious, negative animal welfare

outcomes, lack of full scientific certainty as to the sentience of the animals in question shall not be used as a reason for postponing cost-effective measures to prevent those outcomes” (p. 3).

Birch, following Stephen John, bisects the precautionary principle into two components: what he calls an epistemic rule and a decision rule. The epistemic rule can be formulated as follows:

For the purposes of formulating animal protection legislation, there is sufficient evidence that animals of a particular order are sentient if there is statistically significant evidence, obtained by experiments that meet normal scientific standards, of the presence of at least one credible indicator of sentience in at least one species of that order. (p. 5)

The decision rule is formulated as follows:

We should aim to include within the scope of animal protection legislation all animals for which the evidence of sentience is sufficient, according to the standard of sufficiency outlined in [the epistemic rule]. (p. 5)

Birch provides a great deal of argumentation and clarification to motivate his position, but this outline is clear enough to make some general remarks about the value and unique attributes of the approach. First, as should be obvious, the precautionary approach avoids the most serious pitfalls of precise accounts of sentience since line-drawing occurs not at the point at which we think we have the *best* evidence, but rather at a point where we have *sufficiently reasonable* evidence. So, for example, Birch cites the research program by Robert Elwood that provides evidence that decapod crustaceans engage in motivational tradeoffs involving noxious stimulation. Thus, according to the decision rule provided, “decapods should therefore be brought within the scope of animal protection legislation” (p. 8). Presumably, this would mean, at a minimum, that restrictions be placed on how decapods are killed for food, and methods such as boiling lobsters alive would be prohibited.

The precautionary approach to sentience still has some risk of leaving out sentient animals, in cases where current credible methods for assessing sentience for some reason fail to detect it. But this risk is clearly far lower than that of precise accounts. As such, the chances of causing large amounts of suffering based on a mistaken assessment of sentience are greatly reduced when using precautionary approaches toward policy-making directed at animal protection laws and regulations, as well as determining morally permissible behavior. On the other hand, the risk of mistakenly treating some animals as sentient when they are not increases. Some particular industries, such as fishing, could claim that these economic harms are also serious. But in general, the asymmetry of harms—the potential harm in treating sentient animals as insentient versus the harm of treating insentient animals as sentient—is such that the first type of error carries far more moral risk.

Evidence from the neurosciences has an interesting role to play in precautionary approaches to sentience. It certainly seems conceptually possible that some comparative neuroanatomy evidence, against a background of other scientific knowledge, could count as “statistically significant evidence . . . of the presence of at least one

credible indicator of sentience” [14, p. 5]. As such, neuroscience evidence, in conjunction with behavioral evidence, could likely be sufficient to meet the epistemic condition. But purely neuroscientific evidence by itself probably would not tell us much, since without some understanding of how neural processing leads to behavior, it is difficult to connect this type of evidence to those mental states constitutive of sentience. However, it also seems clear that in some cases the epistemic condition for the precautionary principle could be met without any reference whatsoever to evidence from the neurosciences. Take the “motivational tradeoffs” criteria referenced above in relation to decapods. As it happens, Elwood also has neurological evidence regarding behavior, but one could certainly imagine a set of behaviors sophisticated enough that, even in the absence of any comparative neuroscience, it is sufficient to meet the epistemic criteria. An organism such as a lobster, for example, that typically withdraws from noxious stimulation, but is willing to undergo such stimulation in exchange for particular rewards, and whose tradeoff behavior is altered in the expected ways by the administration of analgesics, would seem to be a good candidate. Thus, though neuroscience evidence can be relevant for precautionary approaches to sentience, it is not strictly speaking required.

Though I think the above suggests that the precautionary principle has an important role to play in determining policy and ethics guidelines, it is also clearly limited in its application to certain types of decisions. For example, in the scenario described at the beginning of the article, we might imagine that the precautionary principle determines that the ants, fish, chickens, and cows all meet our epistemic criteria for sentience. Would it then follow that we should save the ants, even if the evidence for sentience was far stronger in the case of the other animals? It would be unfair to say the precautionary approach as stated above gets the wrong answer in such a case; more accurately, the decision rule above is simply not designed to answer this type of question. The principle is relevant for specific knowledge asymmetries where erring in one direction involves the risk of serious harms. It is not, however, useful for making decisions in situations that require carefully balancing harms and benefits between different options, such as the example at the beginning of this chapter. For that, I turn to a third type of approach to sentience: probabilistic accounts.

13.4 Probabilistic Approaches to Sentience

Returning to the original scenario, the earlier accounts of sentience do not provide plausible guides for decision-making. Relying on a version of a precise account, let us say one’s preferred “guiding theory” of consciousness suggested that a certain type of corticothalamic feedback loop was required for consciousness, thus implying that cows and chickens meet the criteria, but fish and ants do not. This account presumably would suggest that we ought to save the chickens (assuming that cows do not suffer more or enjoy more, though this is not necessarily entailed by the fact that both are sentient). But this approach carries the risk of allowing 35 additional animals to die who could have been saved if your guiding theory was mistaken in

excluding fish. If you think that this is an acceptable risk, what if the choice were between 500 fish and 15 chickens, or between 5000 fish and 15 chickens? Presumably at some point most of us would think the numbers should matter in our decision, but precise accounts that say policy should be determined solely by line-drawing based on the preferred theories would not be sensitive to such concerns.

Similarly, let us say we adapted the precautionary principle such that it suggests that once the epistemic criteria are met, our new decision rule specifies that “all should count for one, and none for more than one” and thus that all sentient lives should be treated equally. But imagine that there were 10 different, plausible evidentiary criteria for sentience and that mammals like cows met all 10, while the chickens only met one. Would we still have an overriding reason to save the 15 chickens over the 10 cows? Again, we can play with the numbers to generate more clearly uncomfortable results. If the choice was between 99 cows and 100 chickens, and we had *dramatically more evidence* that cows are sentient, surely the evidence should be enough to tip the scales at this point. So while the precise accounts are insensitive to numbers of animals, the modified precautionary account is insensitive to strength of evidence, once that evidence has crossed a certain threshold.

As such, we need a different approach. Imagine that we had good reason to think that there is a 95% chance that cows are sentient, an 80% chance that chickens are sentient, a 60% chance that fish are sentient, and a 10% chance that ants are sentient. On one relatively straightforward approach of multiplying the purported value of a particular outcome by the probability of it occurring, we can use those probabilities to assign weightings to different possible choices. Assuming for simplicity that there is no relevant differences in intensity or duration of possible experiences between the animals, and our goal is to maximize the number of sentient lives saved, we could multiply 0.95 by 10 to get an expected difference in the case of cows to 9.5, and multiply the 0.10 by 70 to get an expected difference of 7 for saving the ants. Thus, this approach would favor saving the cows over the ants. However, 15 times 0.8 would get an expected difference of 12 for saving the chickens, and 0.6 multiplied by 50 would get an expected difference of 30 for saving the fish. Thus, this account would recommend saving the fish (on this highly artificial setup) as the approach most likely to save the highest number of sentient animals. Of course, a full ethical accounting would need to take into consideration many other variables, such as how long the animals would live, whether there might be different intensities in emotions, etc., but hopefully the idea of how uncertainty about mental states could theoretically be built into the assessment is clear.

The probabilistic approach to sentience is prominent in certain animal advocacy communities and particularly in what is called the Effective Altruism community. This community is dedicated to doing the *most good*, and as such using probabilities to weight-expected outcomes is a straightforward way of maximizing the expected value of particular actions or policies. Starting from this perspective, it has been argued that it would be better to focus on improving the welfare of chickens over the welfare of cows, since 7 billion broiler chickens are slaughtered for food every year in the USA as opposed to 34 million cows. Though most of these advocates would agree that it is probably slightly more likely that cows are sentient, this slight difference is not enough to make up for the fact that 205 times more chickens are killed. Similar arguments have

been put forward to argue in favor of prioritizing the welfare of fish over chickens and mammals, since their numbers (especially when bycatch is included) dwarf even those of chickens. And, perhaps even more surprisingly, similar numbers-based arguments have been used to suggest that insects should be prioritized over vertebrates even if we think there is a relatively low likelihood of them being sentient.

Of course, such accounts are only as plausible as the idea that we can reliably assign probabilities to the different possibilities. Since we do not truly have access to objective probabilities, the term “credence” is often used to quantify the subjective estimates of probability assigned to particular outcomes. If a new study comes out that says that, say, termites have nociceptors, this presumably should increase the credence we assign to the likelihood that termites are sentient. *How much* it should do so is of course a question that can be answered only in the vaguest terms. Though it is hard to have much confidence in the precise amounts of credence people assign to various bits of evidence, there are presumably some cases, say comparing the welfare of one cow to that of 10,000 chickens, where the differences in numbers are so large as to dwarf the differences that would come from subjective assessments of the likelihood of sentience.

Probabilistic accounts have an important role to play in ethical decision-making. In conjunction with the precautionary accounts, some version of the probabilistic accounts should guide our ethical decision-making in most cases. However, at least as practiced, the application of this framework also has flaws. In particular, it arguably assigns weight to *too many things*. The accounts are sensitive to evidence from any account of sentience that has not been definitively ruled out, which of course includes many mutually exclusive accounts, and thus it stands to reason that some ultimately irrelevant evidence will be playing a role in guiding decisions. Moreover, the approach in practice can (though need not) lead us to adopt an uncritical stance toward different evidence. It can be easy to simply say “person X says fish are insentient therefore I will adjust my credence appropriately” rather than thoroughly and skeptically evaluating the evidence, and it may turn out that faulty reasoning is missed as a result.

13.5 Guiding Ethics and Policy

In light of the above considerations, how ought we go about formulating our ethical and political commitments in light of evidence for sentience in different species? And how relevant are findings from the neurosciences?

In my view, the precautionary and probabilistic accounts are useful for directly guiding these decisions in certain situations, whereas precise accounts are never directly relevant for decision-making, but nevertheless provide important background considerations and guidance. For the sake of ensuring basic protections across nonhuman taxa, precautionary accounts of sentience are needed to prevent the risk of serious harms occurring. Thus, many laws and policies should be formulated that treat animals as sentient provided that there is statistically significant evidence, obtained by experiments that meet normal scientific standards, of the

presence of at least one credible indicator of sentience in at least one species of that order. And the precautionary principle is similarly useful in situations where there is great risk of harm from treating sentient animals as insentient, and relatively little risk of harm from treating insentient animals as sentient.

However, other types of situations involve trade-offs between harms that are more proportionately aligned. In these cases, probabilistic accounts that take into account some rough estimate of the likelihood of sentience, the numbers of animals involved, and an assessment of the magnitude of harms and benefits are more relevant for decision-making. These may not be the only relevant moral criteria in such cases, since concerns about rights, autonomy, and such might be relevant, but it should at least be included as a component in the decision-making process.

However, as noted above, precise accounts must also play an important though indirect role. In particular, without an ongoing search for precise accounts of sentience, we will almost certainly be making incorrect judgments about sentience, and these mistakes have real costs (particularly in cases where we are relying on probabilistic accounts), costs that can include the suffering of sentient organisms and the misappropriation of scarce resources. In searching for precise accounts of sentience, we will often discover that some criteria for sentience are flawed, and this in turn will have implications for how precautionary and probabilistic theories are applied. For example, one commonly suggested criterion for sentience is that an opioid has some effect on behavior in response to a noxious stimulus, with the implication being that if an animal changes their behavior as a result of a painkiller being administered, the original behavior must have been the result of a pain experience. However, we know from neuroscience that there are receptors sensitive to opioids in the peripheral nervous system, and thus reflexive responses to noxious stimulation can be influenced without any changes taking place in the brain. As such, just any behavioral response to opioid analgesics should not, by itself, be regarded as sufficient evidence for sentience, even on a precautionary account. Of course, the principle could possibly be modified such that it does necessarily involve opioids influencing brain systems, and perhaps this modified principle could be used instead. But the upshot is nevertheless that the earlier, more simple rule cannot be justified.

13.6 Conclusion

It is important that the search for precise accounts of sentience continues, and this will certainly involve neuroscience research for the foreseeable future. As connections are made between behavior, self-report, and neural function, criteria for sentience will continue to be refined. This in turn will help to improve precautionary and probabilistic accounts of sentience that are needed to guide policy. Though I think the precautionary principle should help guide the formulation of basic protections, and that evidence from the neurosciences is not *required* in some

cases to conclude that particular species deserve protections, an ongoing search for the neural correlates of sentience must be pursued in order to avoid harms that occur from mistaken accounts.¹

Acknowledgement Thanks to the Alternatives Research Development Foundation (ARDF) and the University of Pennsylvania School of Arts and Sciences for funding the “Animal Research Neuroethics” workshop, which has helped shape my views on these topics, as well as my contribution to the book project more generally. A grant from the Centre for Effective Altruism UK also contributed to the development of the ideas here. The writing of this chapter was supported by Wellcome Trust grant (203132/Z/16/Z). Thanks to the Oxford Uehrio Centre for Practical Ethics, and the Wellcome Centre for Ethics and Humanities for supporting my current work.

References

1. King B. The joys and ethics of insect eating. In: National public radio: cosmos and culture. 2014, April 3. <https://www.npr.org/sections/13.7/2014/04/03/297853835/the-joys-and-ethics-of-insect-eating>. Accessed 2 June 2019.
2. Animal Charity Evaluators. Farmed fish welfare report. 2019. <https://animalcharityevaluators.org/research/other-topics/farmed-fish-welfare-report/>. Accessed 2 June 2019.
3. Broughton J. Descartes’s method of doubt. Princeton, NJ: Princeton University Press; 2003.
4. Metzinger T. Being no one: the self-model theory of subjectivity. Cambridge, MA: MIT Press; 2004.
5. Avramides A. Other minds. London: Routledge; 2000.
6. Degrazia D. Sentience and consciousness as bases for attributing interests and moral status: considering the evidence—and speculating slightly beyond. In: Johnson LS, Fenton A, Shriver AJ, editors. Neuroethics and nonhuman animals. Cham: Springer; 2019.
7. Shriver A. Minding mammals. *Philos Psychol*. 2006;19(4):433–42.
8. Sethu H. How many animals does a vegetarian save? Counting Animals. Accessed at <http://countinganimals.com/how-many-animals-does-a-vegetarian-save/>
9. Varner GE. In nature’s interests? Oxford: Oxford University Press; 2002.
10. Varner GE. Personhood, ethics, and animal cognition: situating animals in Hare’s two level utilitarianism. Oxford: Oxford University Press; 2012.
11. Allen C. Animal pain. *Noûs*. 2004;38(4):617–43.
12. United Nations. Rio Declaration on Environment and Development. Report of the United Nations Conference on Environment and Development, Rio de Janeiro, June 3–14. New York: United Nations. 1992.
13. Steel D. Philosophy and the precautionary principle. Cambridge: Cambridge University Press; 2015.
14. Birch J. Animal sentience and the precautionary principle. *Anim Sent*. 2017;2(16):1.
15. Sebo J. The Moral Problem of Other Minds. *Harv Rev Philos*. 2018;25:51–70.

¹While this paper was in production, I came across a similar article comparing the precautionary principle and the expected utility function by Jeff Sebo. Though the papers have different methodologies and different aims, I highly recommend also taking a look at his insightful paper [15].

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

