

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

François Kammerer
Ruhr-Universität Bochum
kammerer.francois@gmail.com

Ethics without sentience. Facing up to the probable insignificance of phenomenal consciousness^{1 2}

Abstract: Phenomenal consciousness appears to be particularly normatively significant. For this reason, sentience-based conceptions of ethics are widespread. In the field of animal ethics, knowing which animals are sentient appears to be essential to decide the moral status of these animals. I argue that, given that materialism is true of the mind, phenomenal consciousness is probably not particularly normatively significant. We should face up to this probable insignificance of phenomenal consciousness and move towards an ethic without sentience.

Introduction

Phenomenal consciousness appears to be particularly normatively significant. At least some phenomenal states, such as phenomenal pain and pleasure, seem to have a distinctive kind of value – they seem good or bad in virtue of the way they phenomenally feel. This value appears to distinguish them markedly from neighboring non-phenomenal states. For this reason, *sentience* – the capacity to enter phenomenal states – is widely considered the hallmark of a peculiar kind of moral status. Sentient creatures, because they enter phenomenal states endowed with this distinctive sort of value, seem to matter morally in a way in which non-sentient creatures do not. Hence, making progress in the field of animal ethics appears to require primarily an investigation of the distribution of sentience. Which creatures are sentient, and which are not? It seems that we must answer this question if we are to know which creatures matter morally, how much they matter – and how we should treat them.

I argue that sentience should not be seen as the hallmark of a distinctive sort of moral status, because phenomenal consciousness is unlikely to be particularly normatively significant. Indeed, we have good reasons to believe that materialism is true of the mental. If materialism is true of the mental, then it should lead us either to believe that phenomenal consciousness does not exist, or that it exists but is quite different from what it introspectively seems to be. In the first case, phenomenal consciousness,

¹ I want to thank Valentina Martinis and Ismahan Wayah for comments, as well as audience at the Ruhr-Universität Bochum, and two anonymous reviewers from the *Journal of Consciousness Studies*.

² This research was pursued with the support of the Humboldt Foundation.

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

being inexistent, cannot have any normative import in reality. In the second case, we gain two arguments against the view that phenomenal consciousness is particularly significant.

Because phenomenal consciousness and sentience are unlikely to be particularly normatively significant, we should move towards an ethic without sentience. I try to draw lessons from the problems encountered by sentience-based approaches to ethics (and notably to animal ethics) to contribute to the progress of ethics without sentience. At least three such lessons should be drawn. First, taking sentience out of ethics does not necessarily have devastating consequences: it does not threaten the possibility of (animal) ethics. Second, while building an alternative to sentience-based ethics, we should avoid the temptation to find another property that would similarly ground a binary distinction of status across creatures. Third, even if removing sentience from ethics does not necessarily have massive consequences, it would be a mistake to think that it leaves things unchanged.

I first define phenomenal consciousness, and describe its apparent normative significance (§1). Then, I examine the way in which phenomenal consciousness and sentience are often seen as the source of value *par excellence* and the hallmark of moral status. I argue that sentience, in this respect, is used both as a normative magic bullet and as a normative black box (§2). Next, I turn to problems that arise for sentience-based approaches to moral status (§3), before developing my argument for the view that phenomenal consciousness is probably not particularly normatively significant (§4). Finally, I close with lessons for the development of an ethic without sentience (§5).

1. Phenomenal consciousness and its apparent significance

“Phenomenal consciousness” refers to the set of mental states which are such that – to use the now classical expression – there is “something it’s like to be in them” (Nagel, 1974): mental states which are subjectively experienced. Seeing a green patch, feeling a sharp pain in the knee after a long run, smelling freshly grounded coffee: these are supposed to be typical examples of phenomenally conscious states.

Phenomenal consciousness seems to be intrinsically particularly normatively significant. There are many ways to spell out this apparent normative significance. Here is an attempt: it (strongly) seems that the having of phenomenally conscious states – notably *valenced* phenomenal states such as phenomenal pain and pleasure – in itself grounds some significant kind of value – negative or positive. For example, it seems that feeling (phenomenal) pain is *bad* and that feeling pleasure is *good*. At the very least, these feelings seem good and bad *for the subject* who has them. Arguably, it also goes further than that, so

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

that it also seems good or bad *tout court* that a subject feels them.³ Moreover, it seems that phenomenal pain and pleasure are not good or bad merely because they have certain further (good and bad) effects. They seem good or bad *in themselves*, independently of any further effects (feeling pain seems bad in itself, and not just because it creates bad memories). Hence, the value is intrinsic and not merely instrumental.

Besides, this apparent intrinsic value of phenomenal states seems to exist in quite a strong manner – in the sense that the *seeming* itself is strong. In fact, the value of at least *some* phenomenal states – say, the negative value of pain – seems to be extremely obvious, and almost beyond doubt. It seems that we just have to focus on our own phenomenal pains to know directly, without hesitation, that they are *bad*. There appears to be no need for careful reflection on elaborated moral laws, ethical principles, or contractual agreements, to know that *this* (focusing on our current pain) is bad because of the way it *feels*, and hence that, when someone else feels the same thing, it is also bad.⁴ It might even seem that the obviousness of the value of valenced phenomenal states is unique in this respect. The badness of phenomenal pain, for instance, seems particularly obvious and beyond doubt compared to the value of any non-phenomenal events, arguably because phenomenal states “have a subjective component that, so to speak, directly reveals their value” (Seager, 2001, p. 4). This is quite an interesting feature. Indeed, many things can seem *good*, *bad*, or *neutral* at first glance, but we can later come to realize that they are not – that previous impressions and beliefs regarding their value were mistaken. Most Greeks seemed to believe that slavery was not intrinsically bad; Emmanuel Kant thought that masturbation was intrinsically very bad. Yet, I suppose that most of us now think that these judgments were mistaken. Examples of normative disagreement (diachronic and synchronic) abound, and they easily lead to the lingering fear that it could be really hard – if not impossible – to know whether something is *really* good or bad. In contrast, the case of valenced phenomenal states such as pain appears reassuringly simple: of course, pain is bad! There is no doubt to be had there: it is an obvious fact, a fact that we can directly know simply by paying attention introspectively to our pains.

Finally, one notable thing regarding the value of phenomenal states is that this value is supposed to be *particularly significant*, notably compared to the value of neighboring mental states that are *not* phenomenally conscious. Non-phenomenally conscious pain and pleasure, for example (or the non-phenomenal states closest to pain and pleasure, if one likes that better⁵) might very well have *some*

³ I will leave aside here complications that arise when we try to disentangle the various senses of “value” that could be relevant here. A first approximation: “value” can mean both value-for-the-subject, that is, subjective value (the property of situations that are good or bad for someone), and value *tout court*, that is, objective value (the property of situations that are good or bad *simpliciter*). This last sort of objective value can then be understood in different ways – one can distinguish ethical value, aesthetic value, epistemic value, etc. Here the senses that I am interested in are value-for-the subject and objective ethical value.

⁴ (Muehlhauser, 2017) elaborates on this.

⁵ Some might want to reserve “pain” and “pleasure” for phenomenal pain and pleasure.

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

intrinsic value, but they seem, at the very least, to have *significantly* less value than their phenomenal counterparts.

That phenomenal states seem of particular significance compared to neighboring non-phenomenal states manifests itself in the fact that we draw a series of stark normative contrasts. For example, we draw a stark normative contrast between phenomenal and their closest non-phenomenal equivalent. We care a lot about the intense pain that one might phenomenally experience during a medical procedure – arguably, because such pain seems really bad. On the other hand, if, thanks to anesthesia, a patient does not experience phenomenally conscious pain during surgery, their brain might still enter in non-phenomenally conscious states that are the non-phenomenal states closest to phenomenal pain (something like “subliminal pain” or “unconscious pain”) – but we will probably not worry too much. If indeed we fully believe these states to be non-phenomenal – to have no associated subjective experience, “nothing it’s like” to be in them – we will probably judge that they have little intrinsic moral relevance – if at all – and we will not do much to avoid them. They will be a matter of curiosity, not of deep worry.

In a similar fashion, we draw a normative contrast between types of creatures, depending on whether or not we think they are *sentient*; that is, depending on whether or not we think that they can enter in phenomenally conscious states – particularly, valenced states such as pain. At least since Bentham (“the question is not, Can they *reason*? Nor, can they *talk*? But, can they *suffer*?” (Bentham, 1961)) – and probably before him – we take it that we should care much more (if not exclusively) about creatures who (notably) feel pain. We arguably do so because we think that it would be really bad for such a creature to have negatively valenced experiences such as pain. On the other hand, if we are convinced that a given creature is *not* sentient, and cannot phenomenally experience pain, or other negatively valenced phenomenal states (or even cannot experience anything at all), we usually care much less (if at all) about, say, inflicting damage to this creature, or about impacting its mental life in general.⁶

2. Magic bullets and black boxes

Phenomenal consciousness, then, seems particularly normatively significant. Probably for this reason, sentience appears to constitute the ground of a fundamental distinction in moral status amongst

⁶ “Sentience” can be interpreted as referring to the general capacity to enter phenomenal states or the more particular capacity to enter valenced phenomenal states such as phenomenal pain and pleasure. In this paragraph, I did not decide between the two senses. In what follows, I will use the term in the first sense. Although the idea that phenomenal consciousness is particularly normatively significant is more attractive when one focuses on cases of valenced phenomenal states, there are arguments to the effect that even phenomenal states entirely devoid of valence could bear some peculiar kind of value. My argument against the normative significance of phenomenal consciousness will extend to all phenomenal states.

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

creatures. Sentient creatures seem to have a much higher moral status than non-sentient creatures – arguably because their phenomenal states (notably the valenced ones) have a value that is much more significant than the non-phenomenal states had by non-sentient creatures. That sentience grounds a fundamental distinction in moral status is widely accepted and considered obvious by many philosophers and scientists (Birch, 2021; Lin, 2021; Mellor, 2019; Seager, 2001; Sebo, 2018; Shepherd, 2016; Singer, 1993; van der Deijl, 2021). The idea that sentience grounds such distinction in moral status also seems to have been what motivated the “Cambridge Declaration on Consciousness” in 2012 (Low et al., 2012). It is also assumed in the EU Lisbon Treaty and the Global Animal Welfare Strategy of the World Organization for Animal Health (Carruthers, 2019, p. 166; Dawkins, 2021, p. 87). The fact that this idea is so widely accepted has probably to do with the fact that the value of at least some phenomenal states seems itself particularly obvious – as noted above.

The debate in animal ethics thus naturally tends to take the following shape: *sentience* is the property that determines which creatures have a certain kind of particularly significant moral status – and which do not (even if one admits that non-sentient creatures might still have some lower sort of moral status). Are all mammals sentient? Is fish sentient? Are lobsters sentient? What about octopuses, bees, ants, earthworms, etc.? Answering these questions appears crucial because we take the answer to entail a view regarding the moral status of these creatures – and thereby regarding the behavior we should adopt with respect to them. The same could be said of the nascent field of machine ethics (robot ethics, computer ethics): to know whether machines have rights, or dignity, etc. – that is, to know what their moral status exactly is – it seems that we first need to know whether or not they are sentient. Sentience then tends to be seen as some sort of normative “magic bullet” (Carruthers, 2019, p. 166). We see it as a particularly normatively significant property, dispensing a unique higher form of moral status to (and probably only to) the creatures endowed with it, thus grounding a stark normative contrast between the creatures who have it and the other creatures.

Two remarks. First, it is interesting to see how this idea of sentience as a normative magic bullet is widely accepted, while at the same time, there is so much disagreement on what exactly sentience and phenomenal consciousness are and how they are generated/constituted (more on that below). Consciousness researchers are often keen to stress how normatively significant phenomenal consciousness is. It might be because doing so constitutes a convincing manner to obtain research funding (we *really* need to know what consciousness is and where it is located, if we are to know our moral duties!) or more charitably, because it genuinely makes the study of consciousness intellectually, morally – almost existentially – more pressing.

Second, when we consider sentience as a normative magic bullet, while at the same time being ignorant or conflicted regarding what it is and what constitutes/generates it (to say nothing of *how and why* it grounds moral status/value, and *how and why* we can know that it is the case – two questions that

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

contemporary consciousness studies tend to neglect), we treat sentience as some sort of *normative black box*. We assume that it does a certain very-important-normative-job, without knowing much about what it really is or how it does this job. Other entities have been treated as a normative black box in the history of thought. For instance, it is plausible that for many people, the *will of God* was (and still is) a similar sort of normative black box. People agreed that the will of God had to be obeyed, while disagreeing widely about what the will of God was, who could know it and how (to say nothing about why the will of God had to be obeyed – a question that mostly mattered to philosophers). It might also be that *the will of the people* was (and still is) a similar kind of normative black box in the field of politics. The effect of using such a normative black box is *de facto* to replace a normative question (Who has rights, or dignity, or value? What is to be done?) with what seems to be a merely descriptive, non-normative question (Who is sentient? What does God want?), which might appear more tractable (after all, it is indisputably a matter of *facts!*), but also happens to generate intense debate and disagreement in its own right.

3. Problems for a sentience-based approach to animal ethics

Basing animal ethics on sentience leads to difficulties. One of them stems from the fact that there is currently no agreement in the field of consciousness studies regarding the *nature* and – more importantly – the *distribution* of phenomenal consciousness. Different competing theories have widely different implications as to which creatures are sentient (see, for example (Dawkins, 2021, p. 6), for an overview of the different “proposals for membership of the ‘consciousness club’”). This prevents the formation of scientific consensus regarding the related normative issues. To deal with this problem, some researchers advocate a theory-neutral approach to the distribution of animal sentience (Tye, 2017). Some think that, given this state of uncertainty, we should pragmatically set aside – for now – the question of animal sentience when doing animal ethics and base our ethical approach on less controversial considerations – notably, but not only, the satisfaction of animals’ desires – while recognizing that animal sentience could ultimately be highly relevant to the issue at hand (Dawkins, 2021, chapter 8). Others suggest principles to deal with our current state of uncertainty – which should not be equated with a state of complete ignorance – regarding the distribution of sentience, such as precautionary principles (Birch, 2017) or expected value principles (Sebo, 2018). All these considerations recognize that there is a difficulty here, which stems from our relative ignorance – illustrated by the lack of consensus – regarding the nature and distribution of phenomenal consciousness.

However, approaching animal ethics with sentience as a normative black box leads to a much more fundamental problem. It does not stem from our state of ignorance: it is not caused by what we *ignore*,

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

but rather by what we *know* about phenomenal consciousness. Given what we know, we have good reasons to believe that phenomenal consciousness and sentience cannot play the “magic bullet” role that is routinely attributed to them in the normative boxology mentioned earlier. Phenomenal consciousness is (probably) normatively insignificant – in the sense that it is probably *not particularly* normatively significant and cannot justify drawing a stark contrast between states and creatures who are phenomenally conscious and those who are not.

Why so? First of all, I assume that the naturalistic, materialist approach to the mind has been so extraordinarily successful in the last decades that we are justified to conclude that the mind is (metaphysically) nothing over and above a set of “ordinary” material processes (organized in some non-ordinary but also non-intrinsically-mysterious ways). At the fundamental level, there is no special “mental” ingredient: minds are fully constituted by material, i.e., non-mental stuff – substances and properties. I will not argue for this here, as the argument for materialism about the mind has been made elsewhere.⁷ At the same time, there are very strong intuitions and arguments to the effect that phenomenal consciousness is irreducible to anything material (Chalmers, 1996; Jackson, 1982). If we admit all of this, we are essentially left with two main options regarding phenomenal consciousness. Either we are convinced (maybe because of these intuitions and arguments) that phenomenal consciousness cannot be reduced to anything material. In that case, given that we also think that minds are nothing over and above material processes, we conclude that minds are not phenomenally conscious, which naturally leads to (strong) *illusionism*: the view that phenomenal consciousness does not exist, but merely seems to exist (Frankish, 2016; Kammerer, 2021). Or we think that such arguments can be resisted (or, maybe, we find strong illusionism unacceptable?), and we embrace some form or other of (realist) *reductive materialism*, on which every instance of phenomenal consciousness is identical with (or constituted by, or realized by) a set of material processes.⁸ At the same time, it is very difficult to deny that phenomenal states *do not seem* to be such (hence, the strength of the intuitions and argument against materialism). They *do not seem* to be nothing over and above certain material states, which means that our form of reductive materialism will have to include a dose of *weak illusionism* (Frankish, 2016, p. 15-16). That is, it will have to be a view on which, even though phenomenal consciousness exists (and is material), it is *not* like it seems to be – notably, it is not like it seems to be through introspection. Phenomenal consciousness (introspectively) seems to have properties that it does not have in reality. For example, it might seem to be primitive, or immediately known, or irreducibly and qualitatively *feely*, or non-physical, etc. while in reality it is nothing over and above a set of composed

⁷ See for example, among others, (Papineau, 2002, chapter 2).

⁸ Which material processes? The science is not settled. It might be the broadcasting of information in a global workspace, the having of a certain kind of higher-order representation, the activation of certain neural patterns, etc.

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

material states, which cannot be immediately known in any substantive manner, cannot be irreducibly and qualitatively feely, etc.

To recapitulate, I take it that we are justified to believe that either strong illusionism, or a form of reductive materialism that includes a dose of weak illusionism, is true. (Someone who disagrees with this can keep on reading and interpret what follows as a conditional consideration: what would happen if one of these were true?). However, as I will try to show, both of these views have implications when it comes to the normative significance of phenomenal consciousness.

4. Why phenomenal consciousness is probably insignificant

If (strong) illusionism is true, phenomenal consciousness does not exist, which means that nothing is sentient in the strictest sense. Sentience and phenomenal consciousness, then, cannot ground any kind of moral status in reality: being inexistent, they cannot be our normative magic bullet. (Strong illusionists can still admit that there are such things as “quasi-phenomenal consciousness” and “quasi-sentience”, quasi-phenomenal states being non-phenomenal states which are standardly mischaracterized as phenomenal – typically, through introspection. I will explain later why quasi-phenomenal consciousness and quasi-sentience cannot be our normative magic bullet either).

If reductive materialism (including a dose of weak illusionism) is true, then phenomenal consciousness is real, although it is not what it seems to be through introspection. In reality, it is identical with (or constituted by) a set of material processes, even though it appears much different (non-physical, primitive, immediately known, etc.) through introspection. However, in that case, we should not treat it as a normative magic bullet either, for two reasons.

A. The argument from indeterminacy. The first reason is linked to *indeterminacy*. The idea is that this form of reductive materialism, including a dose of illusionism, naturally leads to the view that, in some cases featuring non-human creatures, whether or not phenomenal consciousness is present is *indeterminate*. This, in turn, prevents phenomenal consciousness from playing a key role in grounding moral status.

This sort of reasoning has attracted much attention in recent literature. The idea that plausible forms of reductive materialism lead to the indeterminacy of phenomenal consciousness in many cases, notably for creatures different from adult humans, has been noted and convincingly argued for already quite some time (see notably Papineau, 2002, p. 197-230). The problem it potentially creates for phenomenal consciousness as a ground of moral status, or value, was raised more recently, from different

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

perspectives. (Birch, 2021; Carruthers, 2019; Cutter, 2017).⁹ There are important differences between these lines of reasoning. For instance, Papineau makes his claim for the indeterminacy of phenomenal consciousness by appealing to the “vagueness” of phenomenal concepts, while Cutter appeals to sorite-like cases to argue that reductive materialism implies indeterminacy of phenomenal consciousness.¹⁰ On the other hand, Birch, while drawing inspiration from Papineau’s approach, insists that the sort of indeterminacy that is threatening for consciousness as a ground of moral status does not stem from sorite-like vagueness.¹¹ I think it is safer here to set aside debates surrounding the vagueness of our phenomenal concepts, and focus directly on indeterminacy.¹²

Here is my (brief) version of this concern, mostly inspired by Papineau’s, Carruthers’ and Birch’s work. Given that our introspective grasp of phenomenal consciousness, from this reductive materialist/weak illusionist perspective, is partly illusory, phenomenal consciousness as it exists is *not* the thing that strictly satisfies the conditions borne by our introspective representations. Indeed, there is no such thing. Phenomenal consciousness, then, must be *the real thing out there* that corresponds the best – although it cannot do so entirely – to our introspective representations, typically, because it co-varies as systematically as possible with our introspective representations. The nature of this thing is to be determined by our best scientific theories of the human mind as well as with the help of our favored theory of reference. It could turn out to be a certain kind of broadcasting of information in a global workspace, the having of a certain kind of higher-order representation, or the instantiation of some kind of state described at a neural level, etc. Whatever it is, this thing *is*, or constitutes, or realizes our phenomenal consciousness, and our phenomenal consciousness is nothing over and above it: there is no extra property added on top of that to get to phenomenality.

However, whatever this real thing that is (or constitutes, or realizes) our phenomenal consciousness turns out to be, what seems utterly likely, when we look at different individuals, and even more at different species, is the following. This “real thing out there”, systematically present in our own case when we introspect phenomenal states, will never be exactly replicated in other creatures. For each creature, and even more for each species, there will be differences (sometimes slight, sometimes big) in the kinds of broadcasting of information in global workspaces, or in the kind of higher-order

⁹ A close idea is defended in (G. Lee, 2014). G. Lee does not focus on the indeterminate character of phenomenal consciousness, but he claims that phenomenal consciousness is not naturally significant, and that this endangers its capacity to ground other forms of significance (he focuses on epistemic significance, but it is clear that this could apply to ethical significance).

¹⁰ Cutter uses these considerations to argue *against* reductive materialism (someone’s *modus ponens* is someone else’s *modus tollens*).

¹¹ Moreover, while Birch and Papineau focus on indeterminacy between functional properties and their neural realizers, Carruthers focuses on indeterminacy between different sorts of functional properties.

¹² There are arguments against the view that phenomenal concepts as we have them are vague concepts (Antony, 2006; Simon, 2017), but I do not think that they threaten the view that reductive materialism leads to indeterminacy of phenomenal consciousness – as acknowledged by (Antony, 2006, p. 261) – if we consider (as I do) that materialism is true of the mind, and that plausible forms of reductive materialism require a dose of illusionism – i.e., they require us accepting that phenomenal consciousness is not as it seems to be.

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

representation, etc., that they instantiate. The processes they instantiate will be *similar* in some respects to the processes constituting phenomenal consciousness, but also dissimilar in others; and there will be dissimilarities at various levels of abstractions (from the most abstract – the overall functional structure implemented – to the most concrete – the details of the implementation). Therefore, what these creatures will have is something that *somewhat resembles* (to various degrees) the “real thing out there” present in our case. Will the resemblance be such that the corresponding state also counts as phenomenally conscious, or not – will it be enough for the global broadcasting, the higher-order representation, etc., to be of *the right kind* – the kind that constitutes phenomenal consciousness? It is hard to see how there could be always be a fact of the matter here. We cannot appeal to the conditions borne by our introspective representations to decide this. For non-human creatures, it also seems that we cannot look at what co-varies with *their* introspective representations of phenomenal states (on the very plausible assumption that non-human creatures do not possess them). Hence, the answer will remain indeterminate, in the sense that the truth-value of propositions regarding the phenomenally conscious character of the corresponding creatures will not be settled by facts. At least in some cases, this sort of issue will only be solved by a *semantic decision*. We will have to *decide* whether or not these processes present in other creatures should be called “phenomenal consciousness”.

It is important to understand that this problem is not an epistemic problem: it is not a matter of *ignorance*, as it could be if phenomenality were some extra property that had to be added on top of whatever physical processes co-vary with our acts of phenomenal introspection. Even if we knew all the facts about our own brains and environments and animal’s brains and environments, the problem would remain. Indeed, the problem stems from the fact that the referent of phenomenal consciousness cannot be determined as whatever fulfills a certain defined condition, but as whatever co-varies with a number of introspective acts – which is not enough to yield a determinate answer in all cases.

One might hope that appealing to the right theory of reference (e.g., a view in which natural kinds serve as reference magnets) might help us out of this predicament, as we could claim that our introspective representations of phenomenal consciousness simply refers to the most salient natural property, the instantiation of which co-varies with these representations. Carruthers seems to think that this solution cannot be used in the case of phenomenal concepts, because phenomenal concepts are *not* concepts of natural kinds (Carruthers, 2019, p. 125). However, even if we allow natural kinds to play the role of reference magnets here (for instance, because we think that natural kinds can be reference magnets, even for concepts that are not concepts of natural kinds), I do not think that it is likely to take us out of our predicament. In a nutshell, this is because there will probably be too many salient natural kinds exerting their magnetic powers on our concepts. The choice of the “right” natural kind might then also be a matter of semantic decision. Birch, answering to a comparable line of reasoning he ascribes to Nick Shea (Shea, 2012), gives a – convincing – argument in that direction, by showing that there will be at least two such natural kinds, a neural one and a functional one (Birch, 2021). Another difficulty would come from the

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

fact different contexts of application of introspective representations will plausibly make different kinds of natural kinds salient, and that there will be no non-arbitrary way to decide between those.

If we accept this line of reasoning, this means that, in some cases, faced with the question of whether a given creature is phenomenally conscious, the answer will simply be, not *unknown*, but *indeterminate* – not settled by the facts, but decided merely through semantic decisions, which will have an irreducibly arbitrary aspect. If the presence or absence of phenomenal consciousness is indeterminate in some cases, this means that, in these cases, whether or not a creature is sentient will also be indeterminate.

Now, it seems that value and moral status themselves must be *determinate* properties¹³; or, at the very least, that we cannot be satisfied with a view of value and moral status that treats them as ultimately indeterminate. This is the case, notably, because ascriptions of value and moral status very plausibly play a role in our morally-motivated decisions, and that treating them as indeterminate leads to problematically arbitrary moral decisions – see (Birch, 2021, section 3). Of course, some thinkers – some of those attracted by anti-realist conceptions of ethics – might be willing to admit that value and moral status are, in some sense, indeterminate, so that a full defense of this crucial premise would require much more than these few lines. However, I think that many readers would indeed admit that views on which value and moral status are ultimately indeterminate are unsatisfying. However, admitting this implies that, if sentience itself is indeterminate in many cases, it will not be a satisfying ground of value and moral status. Hence, sentience cannot be particularly significant in a way that would make it able to be our normative magic bullet.

Note that a defender of a non-reductive view of consciousness would not have to face the same consequence. They could maintain that phenomenality corresponds to a range of irreducible – and determinate – phenomenal properties, which are such that there is always a fact of the matter as to whether a given creature instantiates one of these properties – and possesses the corresponding form of moral status – even though we might sometimes be ignorant in this respect.¹⁴

B. The argument from justification. Second, it is not only that reductive materialism/weak illusionism implies that phenomenal consciousness *cannot* play the role of a normative magic bullet – for reasons linked to indeterminacy. It also undercuts the reasons we had to believe that it has the capacity to play

¹³ Some people might want to reject that assumption – on this issue, see, for example, the position Brian Cutter calls “coordinated deflationism” (Cutter, 2017, §7).

¹⁴ There are panpsychist forms of reductive materialism that might also avoid this predicament. I did not focus on them, as they are not among the most discussed by researchers with naturalistic inclinations. The most developed of such views, the Integrated Information Theory (IIT) – which can be interpreted as materialist theory, although it is not always read that way – arguably delivers a determinate, non-stipulative answer regarding the degree of consciousness of every single entity. However, one can note that this theory also implies that some very simple entities – such as the physical realization of an expander graph – have a higher degree of consciousness than a human (Fallon, 2021). If IIT is true, consciousness thus seems very unlikely to constitute the hallmark of moral status. It remains to be seen whether another form of panpsychist reductive materialism could escape the indeterminacy worry and yet provide a view of consciousness that makes it a plausible ground of moral status.

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

this role in the first place – for reasons linked to *justification*. This might seem less threatening: undercutting reasons leave more space than rebutting reasons to defend the undercut belief. However, I take it that this problem is quite serious – maybe even more than the previous one, as it arguably relies on fewer presuppositions. While the concern linked to indeterminacy has been widely discussed in the last years, the concern linked to justification has been mostly overlooked. However, I think it deserves all our attention.

Here is a way to make the significance of the argument from justification salient. The reason why we were so strongly inclined to see sentience as a normative magic bullet in the first place (and then used it as a normative black box) was that the value of some phenomenal states seemed particularly obvious and beyond doubt. While normative skepticism seemed a credible threat in all kinds of non-phenomenal cases, with valenced phenomenal states – most typically, pain – it seemed that we were on sure grounds. Of course, feeling pain is bad – just focus on it and you will see for yourself! So, in spite of persisting ignorance regarding so many aspects of phenomenal consciousness, it seemed that we *knew* that it had this sort of particularly significant intrinsic value that made it able to be our normative magic bullet, because we could introspectively grasp this value in the most secure way.¹⁵ However, if reductive materialism/weak illusionism is true, our introspective grasp of phenomenal consciousness is, to a great extent, illusory: phenomenal consciousness really exists, but it does not exist in the way in which we introspectively grasp and characterize it. This undercuts our reason to believe that certain phenomenal states have a certain value: if introspection of phenomenal states is illusory – if phenomenal states are not as they seem to be – then it means that the conclusions of phenomenal introspection must be treated with great care and a high degree of suspicion, which entails that our introspective grasp of the value of phenomenal states cannot be highly trusted.¹⁶ This does not imply that phenomenal consciousness does

¹⁵ I write as if the value of phenomenal states was grasped through introspection. I think that the real story is somewhat more complicated. For example, in the case of pain, my view is that we grasp the negative value of phenomenal pain *a priori*, simply on the basis of concept application (which is why we can securely *know* that it would be bad to experience a certain kind of pain that we have nevertheless never experienced, as long as we can apply the corresponding concepts) – the relevant concept being the phenomenal concept of pain. However, it is usually through introspection that we form basic phenomenal concepts, and that we come to believe most confidently that some phenomenal properties are instantiated (for the view that *a priori* reasoning is key to knowing the value of phenomenal states, see (Kammerer, 2020, p. 916; A. Lee, 2021)).

¹⁶ I explained in the previous note that in my view, the value of phenomenal states is known *a priori*, while introspection informs us on the instantiation of phenomenal properties – on the real presence of these phenomenal states (note also that I endorse strong illusionism so that, in my view, phenomenal states do not exist – and our *a priori* knowledge of their value therefore is knowledge of non-instantiated properties). Note that the problem would remain the same in this perspective. If we come to know that phenomenal states *as they really are* do not conform to the characterization borne by our phenomenal concepts, our reasons to believe that phenomenal states as they really are have the properties that we were inferring they had (through *a priori* reasoning on the basis of our phenomenal concepts) are also undercut. In fact, the undercutting is even more inevitable in that case: someone who believes that the badness of pain is known directly through introspection could maybe hope to say that, although accepting a dose of illusionism means introspection cannot be fully trusted, we can maintain that it delivers some reliable information, and it could be that information about value is precisely of that kind. (Note that, in that case, the burden of proof will arguably have switched so that defenders of such position will have to find reasons to believe that introspection informs us correctly about value). On the other hand, if one believes that the value of phenomenal states is known *a priori* through reflection on concepts, while introspection informs us

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

not have value, but it undercuts our reason to believe that it does in the first place. Two interpretations are possible here: either we think that this entirely suppresses our justification to believe that phenomenal consciousness has significant intrinsic value, or we conclude that it merely weakens this justification (i.e., the justification we were taking to be overwhelming is in fact weaker, and highly defeasible). However, even in the second option, our impression that phenomenal consciousness has some significant intrinsic value at best becomes on a par with other impressions we might have about other valuable things, which means that we should treat it with great caution and consider it highly defeasible (think about Greeks and slavery, Kant and masturbation, etc.).

Note that this concern arises as strongly even if the proponents of sentience-based ethics manage to avoid the problem of indeterminacy. For instance, even if one identifies a single salient natural kind F which co-varies perfectly with our introspective representations of phenomenal consciousness, and such that F is determinately present *or* absent in every creature we consider (so that each creature is determinately conscious or not), we would still have to wonder why on earth we should think that F brings any particular kind of value. For sure, we were quite sure that F had such value when we introspected it, but we also know now that F is not at all like what we introspected.

Again, defenders of a non-reductive view of consciousness could avoid such consequence. They could claim phenomenal consciousness is exactly what it seems to be through phenomenal introspection. They could also vindicate our high degree of confidence in the belief that phenomenal states have a certain value, by claiming that phenomenality, in virtue of its metaphysical nature, has a unique epistemology (in the strongest interpretation, it is infallibly revealed in introspection).

Combining the two lines of argumentation exposed here, we end up with the following reasoning. If reductive materialism/weak illusionism is correct, then we have reasons to believe that phenomenal consciousness cannot play the role of a normative magic bullet (indeterminacy), while the reasons we had to believe that it does so are undercut (justification). The conclusion is that phenomenal consciousness probably cannot play the role of a normative magic bullet: it is probably normatively insignificant.

Note that the same two lines of argumentation could have been developed, for someone who endorses strong illusionism, about *quasi-phenomenal consciousness* – quasi-phenomenal consciousness being, in a strong illusionist perspective, the (non-phenomenal) “real thing out there” that is standardly mischaracterized, through introspection, as phenomenal. Quasi-phenomenal consciousness would thus also be targeted by the concern linked to indeterminacy, as well as the one linked to justification. This

about the satisfaction of these concepts, as soon as one accepts a substantive dose of illusionism, that means that the characterization borne by the relevant concepts are not satisfied. The reality of any value of any property *a priori* inferred from these concepts becomes fully unjustified, unless one can point at an aspect of the characterization borne by phenomenal concepts that is both *satisfied* by real phenomenal states and *sufficient* to warrant the *a priori* inference to the reality of the value property (something I doubt is feasible).

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

is not surprising, as arguably what strong illusionists call quasi-phenomenal consciousness is nothing but what the reductive materialists/weak illusionists call phenomenal consciousness, the difference being primarily semantic – see (Pereboom, 2019) for an interpretation of the weak illusionist/strong illusionist distinction as a semantic distinction.

To recapitulate: we have reasons to believe that either strong illusionism, or a weak illusionist form of reductive materialism, is true of phenomenal consciousness. If any of these two is true, we should abandon the idea that phenomenal consciousness is a normative magic bullet. That is the case, either (a) because being non-existent, phenomenal consciousness cannot play this role – it cannot have any normative significance in reality. Or, (b) because phenomenal consciousness (and the same would apply to quasi-phenomenal consciousness if we thought phenomenal consciousness to be non-existent) is existent, but is such that we have reasons (from indeterminacy) to believe that it cannot play that role, as well as reasons (from justification) to consider our belief that it plays this role as undercut. We should conclude that phenomenal consciousness – and therefore sentience – is probably normatively insignificant.

We can compare this with the case of God. To someone who considers the will of God as a normative magic bullet, many naturalistically minded philosophers would object the following. From a naturalistic perspective, either God does not exist, or it exists but is identical with some natural entity (maybe the whole of nature, or some hypothetical aliens running the simulation in which we live, etc.). In the first case, it cannot play any normative role in reality. In the second case, (a) our antecedent reasons (appealing, say, to his absolute Benevolence) to believe that He should be obeyed are undercut; (b) we also have direct reasons to think that God, as it exists, cannot fill the bill anyway – that it is simply not the kind of entity able to give the right kind of normative glow to possible courses of actions. I take it that we are in a comparable situation regarding phenomenal consciousness.

5. Ethics without sentience

We now need to face up to the fact that phenomenal consciousness is probably normatively insignificant. This should lead us to revise our conception of ethics and move towards an ethic without sentience – whether we focus on human ethics, animal ethics or machine ethics. What will an ethic without sentience look like? I cannot do much more here than merely scratch the surface of the issue, but I will indicate what I take to be some major elements to keep in mind when dealing with it – mostly with animal ethics in mind.

One first thing to keep in mind is that moving towards ethics without sentience does not necessarily require us to proceed to drastic changes in our approach to ethics – when it comes to our broad

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

understanding of which situations are valuable and which creatures have moral status. It certainly does not mean that the very possibility of ethics – or of animal ethics – is threatened. We can easily recognize that phenomenal consciousness is normatively insignificant without concluding that *nothing* is significant. Moreover, we can recognize the insignificance of phenomenal consciousness, while maintaining that many of the situations we deemed valuable (and many of the creatures we judged to have moral status) really are valuable (and really have moral status), even if it turns out that this value (and this status) have nothing to do with phenomenality. There are numerous coherent and attractive approaches to ethics – applicable to animal ethics – that do not appeal to phenomenality. For example, one can maintain that what makes certain valenced mental states good or bad is not how they phenomenally feel, but the fact that they are partly constituted by the satisfaction or the frustration of a conative state – typically, a desire or a preference (Carruthers, 2019, p. 171-174; Dawkins, 2021, chapter 2).¹⁷ When an animal undergoes pain, this is bad, not because of the way pain phenomenally feels, but because, at least in standard cases, the animal *wants it to stop* – sometimes very intensely – and this desire is frustrated. Such a desire-based approach to the value of mental states – of humans and animals – seems to escape the problems plaguing the consciousness-based approach because our concept of desire (and our apprehension of what makes the satisfaction of desires valuable) does not appear essentially tied to introspection and to the first-person, but rather seems analyzable in functional, third-personal terms. Hence, materialism about the mind does not raise the concerns linked to *indeterminacy* and *justification* in the case of desires – at least not in the same way as it does for phenomenal consciousness. The concern regarding indeterminacy does not arise in the same way because our concepts of desire are arguably analyzable in functional terms. Therefore, we can in principle identify desires with whatever fulfills the relevant functional role – whatever satisfies the characterization borne by our concept of desire – and thus in principle determine which creatures have them. (We might remain relatively ignorant in this respect, but there will be a fact of the matter as to whether a creature has certain desires). In contrast, we could not do that with phenomenal consciousness, because, as shown earlier in the argument from indeterminacy, accepting a dose of illusionism leads us to admit that phenomenal consciousness cannot be identified with whatever satisfies the characterization borne by our introspective representations (to which our concept of phenomenal consciousness are essentially tied). The concern regarding justification does not arise either, because, if we make sure that the concept of desire we use is pure of phenomenal components, whatever reasons we have to believe that desires ground value will be left untouched by the suspicions regarding phenomenal introspection brought by materialism.¹⁸ Investigation of animals' desires, moreover, certainly seems more amenable than phenomenal consciousness to study by a variety of standard third-person methods (see Dawkins, 2021

¹⁷ Carruthers thinks that consciousness is insignificant for reasons linked to indeterminacy. Dawkins does not think so, but she argues that we should (pragmatically) approach animal ethics without reference to phenomenal consciousness and sentience because these notions are too contested to be useful in that respect.

¹⁸ Desire-based approaches to animal ethics are not the only possible replacement for sentience-based ones. For example, one could also opt for a reason-based Kantian approach.

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

for a methodological overview). This is why a researcher like Dawkins thinks that we should ground our conception of animal welfare on animals' desires (among other things), as long as we have not made sufficient progress on the study of phenomenal consciousness. I concur with this recommendation, except that, for the reasons given above, I think that this switching away from sentience and phenomenal consciousness in animal ethics should be definitive.

The second thing to keep in mind is that, even if we devise a satisfying alternative ethical approach, it might nevertheless be wise *not* to expect it to provide an alternative normative magic bullet. It might be wise *not* to expect our ethical approach to describe a set of properties such that they – and only they – have a certain kind of distinctive normative significance, in virtue of which they endow certain creatures with a peculiar moral status, and so that they allow drawing a stark all-or-nothing contrast between mental states and creatures who have these properties and those who do not. Indeed, even if I do not think that materialism implies the same dose of illusionism about (for instance) desires than about phenomenal consciousness, it is nevertheless the case that our ethical conceptions are first going to be formulated and justified at the level of folk (psychological) concepts, featuring a relatively small number of distinctions. The scientific study of the corresponding mental states, on the other hand, reveals much more variety, complexity, and degrees. As noted above, this does not create the problems faced by sentience-based ethics, as it does not necessarily lead to indeterminacy: if even a simple concept of “desire” is defined functionally, then there is no indeterminacy in whether or not a certain creature desires anything. (The concept simply ends up carving out a simple, binary distinction where reality harbors much, much more psychologically relevant differences). However, it prevents our ethical conceptions from being substantively enriched by scientific discoveries. It seems much more reasonable to formulate our ethical guidelines in a way that makes the corresponding concepts open to extensions and refinements. I also think that, at least in the case of desire-based ethics, this is considerably more intuitive. Arguably, the reason why we intuitively judge that desire-satisfaction is normatively significant is that when a creature desires something, this thing starts *matter*ing to them. Intuitively, we also think that there are many ways, with many possible degrees of sophistication, in which something can matter to a creature, and that these ways and these degrees have themselves normative import. On the one extreme, we can think of the most basic way of desiring: a creature can value negatively or positively certain state of affairs, grasped in the roughest way through some basic sensing system. On some views, entities as simple as bacteria can do that (Lyon & Kuchling, 2021). On the other hand, we can think of the most sophisticated ways of desiring. Creatures such as, at least, humans, can desire for a thing to thrive in what they take to be its own proper way to thrive *and* at the same time desire their own desire for this thing to thrive to persist – an attitude close to what Harry Frankfurt called “caring” (Frankfurt, 1988). Between the two, we intuitively admit that there is some kind of progressive and multidimensional scale of desires, which is normatively relevant – states of caring matter more than the most basic desires. When moving towards an ethic without sentience, we would be wise to ground our

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

ethical system on concepts that we will treat as complex and degreed, and even more as “complexifiable” as the study of human, animal and artificial minds progresses.

A third point to keep in mind is that, while we should not think that taking sentience out of (animal) ethics necessarily leads to devastating consequences, it would also be a mistake to conclude that it would be entirely inconsequential. Peter Carruthers, who argues at length for the view that phenomenal consciousness should not be used as a normative magic bullet in animal ethics for reasons linked to indeterminacy, concludes that “much can remain unchanged. [...] So anyone who thinks, for whatever reason, that one is obliged to prevent bad things from happening to animals, can continue with their views unchanged” (Carruthers, 2019, p. 177). However, if we take seriously the concern linked to justification, we have to admit that “whatever reason” we had to care about certain mental episodes in humans and animals is *not* left untouched. Part of the reason we had to care about, say, animal pain, was that we had an introspective grasp of the intrinsic badness of phenomenal pain in our own case – that we took to be quite certain – from which we concluded that (phenomenal) pain would then also be intrinsically bad in the animal case. However, if I am correct, we should *not* take our grasp of the intrinsic badness of our (phenomenal) pain to be particularly well-grounded – if at all (for an exploration, from a Nietzschean perspective, of the view that pain might not be intrinsically bad, justified by doubts on introspection, see (Delon, forthcoming)). We cannot have our cake and eat it: the same views that lead us to admit that phenomenal consciousness is probably normatively insignificant, and unable to play the role of a normative magic bullet, also undercut some of our most crucial ordinary way of justifying our beliefs regarding the value of various mental episodes, which appeals to their phenomenal feel. We cannot tacitly rely on these un-interrogated beliefs while we move towards an ethic without sentience.

Determining what consequences will follow exactly from us taking the problem of justification seriously falls beyond the scope of this article. It obviously depends on many factors – notably, on the determination of the available alternative ways of justifying normative beliefs. One thing that seems useful, however, when moving towards an ethic without sentience, is an understanding of how and why we came to see phenomenal consciousness’ value as something that could be known in a uniquely safe way – in contrast with the value of other things. I suspect that something like a *normative relocation story* might be correct. The standard “relocation story” about sensory qualities says approximately the following: while we first intuitively located secondary sensory quality (redness, greenness, softness, etc.) in external objects, since the Galilean revolution and the mathematization of physics, we were led to believe that these qualities could not literally inhere in the external, physical world. We relocated them *in ourselves* (typically, in our phenomenal experiences). In the realm of phenomenal consciousness, things could be as they appear to be, redness and greenness could exist in all their qualitative glory, and they could be given to us in the most direct manner. I think that a similar origin story might convincingly be told about value. The resulting *normative relocation story* says that we first intuitively located *values* (goodness, badness) in *external objects and events*: we saw the fresh beverage

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

we desired as *good* in itself, the threatening predator we feared as *bad* in itself. However, we started to believe (long before Galileo, and at least since Protagoras, Democritus and Epicurus) that such values could not really inhere objectively in external objects and events. We concluded that they could only exist in our experiences (typically, our pains and pleasures), where they could be what they appear to be and be given to us in the most direct manner.

Materialism about the mind forces us to abandon the view that sensory qualities really exist in the phenomenal realm as they appear to be. I am inclined to conclude that it also leads us to draw a similar conclusion about value. Which satisfying ways to justify our beliefs about value subsist once we cease to locate value in the phenomenal realm and abandon the idea that this phenomenally-located value is introspectively knowable with great certainty? How much our concept of value itself must therefore be revised? These are important questions that will have to be dealt with to make progress with our ethic without sentience.

6. Conclusion

We should move to an ethic without sentience because sentience is unlikely to have the kind of peculiar normative significance required by sentience-based approaches to ethics. This does not threaten the possibility of ethics, but it would be a mistake to think that this move could be entirely inconsequential.

7. References

- Antony, M. (2006). Vagueness and the Metaphysics of Consciousness. *Philosophical Studies*, 128(3), 515-538.
- Bentham, J. (1961). *An Introduction to the Principles of Morals and Legislation*. Doubleday.
- Birch, J. (2017). Animal Sentience and the Precautionary Principle. *Animal Sentience*, 2, 16(1).
- Birch, J. (2021). Materialism and the moral status of animals. *Manuscript*.
<https://philpapers.org/rec/BIREAT-4>
- Carruthers, P. (2019). *Human and animal minds : the consciousness questions laid to rest*. Oxford University Press.

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

Chalmers, D. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.

Cutter, B. (2017). The Metaphysical Implications of the Moral Significance of Consciousness. *Philosophical Perspectives*, 31(1), 103-130.

Dawkins, M. S. (2021). *The Science of Animal Welfare: Understanding What Animals Want*. Oxford University Press.

Delon, N. (forthcoming). Strangers to ourselves: a Nietzschean challenge to the badness of suffering. *Inquiry: An Interdisciplinary Journal of Philosophy*.

Fallon, F. (2021). Integrated Information Theory of Consciousness. In *Internet Encyclopedia of Philosophy*. <https://iep.utm.edu/int-info/>

Frankfurt, H. G. (1988). *The Importance of What We Care About: Philosophical Essays* (1st ed.). Cambridge University Press.

Frankish, K. (2016). Illusionism as a Theory of Consciousness. *Journal of Consciousness Studies*, 23(11-12), 11-39.

Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32(April), 127-136.

Kammerer, F. (2020). The Normative Challenge for Illusionist Views of Consciousness. *Ergo*, 6(32), 891-924.

Kammerer, F. (2021). The illusion of conscious experience. *Synthese*, 198, 845-866.
<https://doi.org/10.1007/s11229-018-02071-y>

Lee, A. (2021). How We Know Pain Is Bad. *Manuscript*. https://ec60f7fc-41a3-4918-9f3d-6f5b65c3308a.filesusr.com/ugd/2dfbfe_f4b9fb30a8e141449091f26bafa33196.pdf

Lee, G. (2014). Materialism and the Epistemic Significance of Consciousness. In U. Kriegel (Ed.), *Current Controversies in Philosophy of Mind* (p. 222-245). Routledge.

Lin, E. (2021). The experience requirement on well-being. *Philosophical Studies*, 178(3), 867-886.

Low, P., Panksepp, J., Reiss, D., Edelman, D., Van Swinderen, B., & Koch, C. (2012). *The Cambridge Declaration on Consciousness*.
<https://fcmconference.org/img/CambridgeDeclarationOnConsciousness.pdf>

This is a draft. The final version of this article is forthcoming in the *Journal of Consciousness Studies*. Please refer to the final published version.

- Lyon, P., & Kuchling, F. (2021). Valuing what happens: a biogenic approach to valence and (potentially) affect. *Philosophical Transactions of the Royal Society B*.
<https://doi.org/10.1098/rstb.2019.0752>
- Mellor, D. J. (2019). Welfare-aligned Sentience: Enhanced Capacities to Experience, Interact, Anticipate, Choose and Survive. *Animals*, 9(7), 440.
- Muehlhauser, L. (2017). Report on Consciousness and Moral Patienthood. *Open Philanthropy*.
<https://www.openphilanthropy.org/2017-report-consciousness-and-moral-patienthood>
- Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83(October), 435-450.
- Papineau, D. (2002). *Thinking about Consciousness*. Oxford University Press.
- Pereboom, D. (2019). Russellian Monism, Introspective Inaccuracy, and the Illusion Meta-Problem of Consciousness. *Journal of Consciousness Studies*, 26(9-10), 182-193.
- Seager, W. (2001). Consciousness, Value and Functionalism. *Psyche*, 7, 1-19.
- Sebo, J. (2018). The Moral Problem of Other Minds. *The Harvard Review of Philosophy*, 25, 51-70.
- Shea, N. (2012). Methodological Encounters with the Phenomenal Kind. *Philosophy and Phenomenological Research*, 84(2), 307-344.
- Shepherd, J. (2016). Moral Conflict in the Minimally Conscious State. In W. Sinnott-Armstrong (Ed.), *Finding Consciousness: The Neuroscience, Ethics, and Law of Severe Brain Damage*. Oxford University Press.
- Simon, J. A. (2017). Vagueness and zombies: why 'phenomenally conscious' has no borderline cases. *Philosophical Studies*, 174(8), 2105-2123.
- Singer, P. (1993). *Practical Ethics*. Cambridge University Press.
- Tye, M. (2017). *Tense Bees and Shell-Shocked Crabs: Are Animals Conscious?* Oxford University Press.
- van der Deijl, W. (2021). The sentience argument for experientialism about welfare. *Philosophical Studies*, 178(1), 187-208.