

DOI: 10.1017/psa.2022.59

This is a manuscript accepted for publication in *Philosophy of Science*.

This version may be subject to change during the production process.

Should animal welfare be defined in terms of consciousness?

Jonathan Birch

Centre for Philosophy of Natural and Social Science,
London School of Economics and Political Science,
Houghton Street, London, WC2A 2AE, UK.

j.birch2@lse.ac.uk

<http://personal.lse.ac.uk/birchj1>

Abstract

Definitions of animal welfare often invoke consciousness or sentience. Marian Stamp Dawkins has argued that to define animal welfare this way is a mistake. On Dawkins's alternative view, an animal with good welfare is one that is healthy and "has what it wants". The dispute highlights a source of strain on the concept of animal welfare: consciousness-involving definitions are better able to capture the normative significance of welfare, whereas consciousness-free definitions facilitate the validation of welfare indicators. I reflect on how the field should respond to this strain, ultimately recommending against splitting the concept and in favour of consciousness-involving definitions.

1. Animal consciousness and welfare

What is the relation between consciousness and animal welfare? Let us note from the outset that there are various things one might mean by “consciousness”. My main focus here is subjective experience, or what philosophers tend to call “phenomenal consciousness”. To say that a state of an animal is phenomenally conscious is to say that there’s something it feels like to be in that state (Block 1995). Ideally we would have a better definition than this, but phenomenal consciousness a notoriously difficult-to-define property, one that resists definition in explicitly functional terms. What we can do is point to examples (Schwitzgebel 2016). When you’re in state of dreamless asleep, in a coma or under general anaesthetic, there’s nothing it feels like to be you. But as we go through our waking lives, there is something it feels like to experience odours, tastes, colours, sounds, pleasures, pains. This is the sense of “consciousness” that gives rise to the notorious “hard problem of consciousness”, the problem of trying to explain how neural activity could give rise to states that feel like something (Chalmers 1995).

Animal welfare scientists often take it for granted that consciousness in this sense is intimately related to animal welfare. Definitions of animal welfare typically appeal to sentience, consciousness, experience, subjective feeling, or related ideas. These definitions matter in so far as they shape the aims of the field. For example, Donald Broom (2008), defines welfare as an animal’s “state as regards its attempts to cope with its environment” but adds that “welfare includes health and the extent of positive and negative feelings”. While Broom considers subjective feelings “an extremely important part” of welfare, Ian Duncan (1993) has argued for a definition of welfare that considers *only* subjective feelings: “Neither health nor lack of stress nor fitness is necessary and/or sufficient to conclude that an animal has good welfare. Welfare is dependent on what animals feel” (Duncan 1993).

In the 1980s, Marian Stamp Dawkins was a pioneer of this way of thinking about animal welfare, writing that “to be concerned about animal welfare is to be concerned with the subjective feelings of animals, particularly the unpleasant subjective feelings of suffering and pain” (Dawkins 1988). More recently, however, Dawkins (2012, 2021) has argued for the importance of defining welfare *without* appealing to consciousness or related ideas. Dawkins now favours a view on which animal with good welfare is one that is healthy and “has what it wants” (2021, 11). I aim to reconstruct charitably what I see as Dawkins’s strongest argument against consciousness-involving definitions of welfare, before turning to Dawkins’s alternative and its problems. I reflect on the strain on the concept of animal welfare the debate reveals, and consider how the field of animal science should respond.

2. Two unpersuasive considerations

Dawkins has offered various different reasons from stripping the concept of consciousness out of the concept of welfare, not all of them persuasive. Perhaps the least persuasive is a direct appeal to the hard problem, as if the hard problem immediately led to scepticism about animal consciousness:

Even with our own consciousness, we still do not understand how the lump of nervous tissue that makes up our brain gives rise to private subjective experiences [...]. And because we do not understand how the human brain makes us conscious, we do not know what to look for in other species to decide if they, too, have

conscious experiences like us. Perhaps they do, but how would we know? (Dawkins 2021, 5)

This is too quick. If evidence-based attributions of consciousness relied on solving the hard problem, we would not be able to make evidence-based attributions of consciousness to other humans. While there may be limits to the level of certainty we can achieve in the human case, human verbal reports of experiences provide good abductive evidence—if they did not, there could be no science of human consciousness. The key question is whether non-verbal indicators provide enough evidence to get a parallel science of animal consciousness off the ground. The mere existence of the hard problem does not imply that they cannot.

Also unpersuasive is an appeal to substantial disagreement in the science of human consciousness. As Dawkins notes, there is great disagreement about the correct theory of consciousness, with currently fashionable options including the global workspace theory (Mashour et al. 2020) and the integrated information theory (Oizumi et al. 2014) and great disagreement about the neural basis of human consciousness too, with proponents of different theories tending to favour different brain areas (Boly et al. 2017; Mashour et al. 2020). Yet substantial disagreement about the nature of a property is not a compelling reason to deny that the property is a constituent of welfare. If disagreement were itself a reason, we would also need to avoid defining welfare in terms of preferences or health, since there is plenty of disagreement about the nature of these properties too.

3. The challenge from unconscious affect

Dawkins, however, also has an empirical motivation for defining welfare in a consciousness-free way, and I see this as the strongest part of her case. Dawkins (2021, 9-11) points to a body of work on human affect in which conscious and unconscious motivational states are apparently dissociated (LeDoux and Pine 2016).¹ In a typical experiment of this genre, an emotionally salient stimulus (such as an angry or happy face) will be presented to subjects in two ways: subliminally (i.e. below the subjective threshold of conscious perception) and supraliminally (i.e. above that threshold). It turns out that some arousal responses interpreted as “fear responses”, such as breaking out in a sweat (as measured by skin conductance), are evoked by the subliminal stimulus. When the stimulus is supraliminal, these responses tend to be weaker, perhaps indicating top-down inhibition of arousal (Tamietto & de Gelder 2010). Moreover, in cases of “affective blindsight”, a subject is unable to consciously perceive the presented face due to damage to their visual cortex but is nonetheless able to guess its valence (e.g. happy or angry) better than chance. Evidence of this type has led Joseph LeDoux and collaborators to posit a “two-system” model of affective processing, in which subjective feelings depend on a cortical pathway, whereas “defensive survival circuits” can be activated by a subcortical, entirely unconscious pathway (LeDoux & Pine 2016; LeDoux & Brown 2018).

The evidence for the two-system model is suggestive but not overwhelming. Granting that the stimuli were indeed unconsciously perceived (cf. Peters et al. 2017), a remaining methodological challenge is to show that the *entire pathway* from stimulus to defensive response, and not just the perceptual part of it, occurred non-consciously. The possibility of unconscious vision eliciting

¹ Paul et al. (2020) have also emphasized the challenge to animal welfare science posed by this literature, but without advocating a shift to a consciousness-free definition of welfare.

amygdala activity which is *experienced as conscious affect* before producing behavioural consequences is hard to rule out.

Indeed, there is some support for this latter possibility from cases in which subjects with affective blindsight have been able to *verbally report* the valence of the feeling evoked by an unseen stimulus in addition to guessing correctly in forced-choice tasks (Anders et al. 2004). There is also some evidence pointing the other way from a study in which some participants (presented with subliminal faces) were instructed to use their feelings as a guide to the valence of the face and yet failed to perform any better than controls (Bornemann et al. 2012). However, although it is important to LeDoux and colleagues' case, the evidence from this experiment is weak, since only 4/19 subjects in the group instructed to use their feelings actually attempted to do so, according to a post-experiment questionnaire (Bornemann et al. 2012, 121).

The evidence for the two-system view of affect need not be overwhelming or even strong, however, to present animal welfare scientists with a challenge. We need only grant that it is *credible*. To the extent that it is credible, then it is credible that many defensive responses to stimuli in animals do not indicate conscious affect. We should take seriously the possibility that such responses in non-human animals are controlled by an analogue of a subcortical pathway that is non-conscious in humans. Intuitively, the idea that the conscious pathway could be absent altogether deserves increasingly serious consideration as the evolutionary distance from humans increases. Even among mammals there are some differences of neocortical organization, but non-mammals do not possess a neocortex at all, so they either have no conscious pathway or else have evolved a different neural route to the same result.

I see this as the basis for Dawkins's most persuasive argument. Without reliable markers of conscious as opposed to unconscious affect, or a reason to consider unconscious affect unlikely, we have no way of confidently validating putative indicators of subjective feelings. We can reconstruct the argument like this:

The challenge from unconscious affect:

1. Animal welfare should be defined in such a way as to make it feasible (with current methods) to validate animal welfare indicators with high confidence.
2. We do not currently have reliable methods for distinguishing conscious from unconscious affect in animals.
3. If we define animal welfare in terms of conscious affect, then, in the absence of reliable methods for distinguishing conscious from unconscious affect, we will be unable to validate animal welfare indicators with high confidence.
4. Animal welfare should not be defined in terms of conscious affect.

To give one example: a currently popular indicator of animal welfare is judgement bias, where an animal is presented with an ambiguous stimulus and its response is taken to indicate an "optimistic" or "pessimistic" expectation of reward (Mendl et al. 2009). Optimistic expectations are taken as a sign of good welfare, whereas pessimistic expectations are taken as a sign of poor welfare. Dawkins's basic point is: for all the talk of "pessimism" and "optimism", we have no reason to think this indicator tracks subjective feelings rather than unconscious affective states, and no

reliable way of finding out (Dawkins 2012, 109-110). So, if welfare constitutively involves subjective feelings, such an indicator cannot be confidently validated.

4. What's the alternative?

Suppose we accept this argument. What is the alternative? Dawkins (2021) proposes that we define animal welfare in terms of health and “valenced states”, where the concept of a “valenced state” is characterized functionally in a way that is neutral between conscious and non-conscious processing.

The idea is that all animals have at least a minimal capacity for preference. Presented with options, their behaviour will often show a systematic bias. Moreover, there will typically be some capacity for reinforcement learning: some states will lead to positive reinforcement and others to negative reinforcement. Although this departs from current usage, we could stipulate that the terms “positive and negative valence” are to refer to the functional significance of a state for preference and reinforcement learning without saying anything about whether or not the state is subjectively felt. Then we can say that good welfare for an animal consists in its (a) being healthy and (b) having many positively valenced states and few negatively valenced states. The best balance of positive and negative valence is open for debate, since negative valence has value for learning. I take this to be the idea that is captured in Dawkins’s slogan: an animal has good welfare if it is healthy and “has what it wants”.

Intuitive motivation for the view comes from cases such as a bird that flies continually into the bars of its cage (Dawkins 2021, 8, 49). Dawkins argues: it is obvious that the bird has compromised welfare, but it is not obvious that the bird is experiencing conscious affect. The proposed account of welfare can make sense of this, because it *is* obvious that the bird is in a state that motivates escape behaviour. The intuition is: we should not want a science of animal welfare that makes it difficult, perhaps even impossible, to assess confidently whether the bird’s welfare has been compromised.

5. The normative challenge to Dawkins’s alternative

Dawkins’s alternative has a significant downside: it does not fulfil the *normative* role of the concept of animal welfare. It fails what we might call the criterion of “normative aptness”. Animal welfare science aims to achieve normatively significant results. It is what David Fraser (2008, following Liora Salter 1988) has called a “mandated science”: a science with a societal mandate to inform animal welfare policies and to improve welfare. I contend that, if animal welfare science is to fulfil this mandate, then the construct of animal welfare should denote a property that deserves high ethical priority and is a genuine source of ethical constraint on our treatment of animals.²

To assess whether a construct can play this role, we cannot avoid substantive ethical questions. We cannot place a firewall between the science and ethics of animal welfare. We must ask which properties are the source of ethical constraints. I contend that health and valence (in the above sense) are not enough: they miss the distinctive importance of *experienced valence*. If a valenced state is not experienced by the animal, then it either has no ethical significance at all or else it has far less significance than an experienced state. This is implied by ethical theories that give a central

² Closely parallel claims can be made about the construct of human wellbeing. See Alexandrova (2012).

role to subjective experience (such as hedonic utilitarianism and Korsgaard’s (2018) neo-Kantian view) but I take it to be a principle that is more plausible than any particular theory.

As a motivating example, consider *Paramecia aurelia*, a single-celled ciliate. A paramecium has no nervous system, but it does have health, and it has states that drive defensive/escape responses—valenced states, in Dawkins’s stripped-back sense. As Michel (2019) has noted, “if a paramecium encounters a potentially dangerous concentrated salt solution or acetic acid, it will back away and swim in a different direction or engage in defensive behavior by discharging trichocysts” (Michel 2019, p. 2418). Yet failing to consider these states is either not an ethical failing or, at worst, it is a failing far less grave than failing to consider the states of a conscious animal. The presence or absence of consciousness makes a large ethical difference.

If consciousness were ethically significant in theory yet irrelevant to current animal welfare policy, it might be tempting to dismiss the practical importance of its omission. But consider the question of whether to protect any invertebrate animals—and, if so, which ones. Cephalopod molluscs are protected in science in the EU and UK but not in any non-scientific context. Switzerland recently brought in legislation to ban the live boiling of decapod crustaceans, and a recent report to the UK government recommended the same move (Birch et al. 2021). It is part of the mandate of animal welfare science to inform policy decisions like these by providing evidence about which slaughter and processing methods create high risks of experienced suffering. There is no escaping the entanglement of these questions with questions of consciousness.

In short, Dawkins’s proposed definition sacrifices normative aptness to enable confident validation of welfare indicators with current methods, whereas definitions that emphasize subjective feeling opt for normative aptness at the cost of making welfare indicators challenging to validate. We would ideally like a definition that is normatively apt *and* reliably tracked by well-validated indicators, but in fact these two considerations seem to pull in opposite directions, at least given our current state of knowledge (**Figure 1**). Moreover, this is not merely a semantic dispute—it is a clash between two ways of thinking about the aims and priorities of animal welfare science.

□




Figure 1: Strain on the concept of animal welfare

6. Benefits and risks of conceptual splitting

How to resolve such a dispute? I think it may help to reflect on an analogy with the concept of consciousness itself, and its role in human consciousness science. Unlike animal welfare science, human consciousness science is *not* a mandated science; it does not generally aim to inform policy decisions. But it does face a comparable problem, in that disputes about the meaning of its central term reflect bigger-picture disputes over the aims and priorities of the field.

Ned Block (1995) highlighted one persistent source of strain. To study consciousness, scientists rely on the idea that verbal reports and other voluntary actions will reliably track conscious states. This pulls in the direction of defining consciousness in a way that would secure this link with voluntary report. Yet what brings scientists to the topic of consciousness in the first place is often a drive to understand “what it’s like” to be a conscious subject—a drive to address the hard problem. This pulls in the direction of defining consciousness in a way that puts the hard problem and “what it’s like” centre-stage, avoiding any *a priori* commitments regarding functions.

Block’s solution was to split the concept of consciousness (**Figure 2**). “Access consciousness” (sometimes also called “conscious access”) refers to the availability of information for planning, reasoning and voluntary action, whereas “phenomenal consciousness” refers to “what it’s like”, subjectively, to be you. Block claimed to be analysing the ordinary concept, but the move is often viewed in retrospect as a piece of conceptual *engineering* rather than analysis (Chalmers 2020). The distinction is now widespread, but not universal, in the field. A similar splitting move is possible for animal welfare. We could take Dawkins’s concept and relabel it “functional welfare”, while relabeling a purely experience-based definition such as Duncan’s as “experienced welfare” (**Figure 3**).




Figure 2: Strain on the concept of consciousness, leading to its split in the 1990s into “access consciousness” and “phenomenal consciousness”.



Figure 3: A possible response to the strain depicted in Figure 1, in which we split “animal welfare” into “functional welfare” and “experienced welfare”.

But should we? I think we could expect a similar profile of risks and benefits to those of the access/phenomenal split. In consciousness science, the split has facilitated the development of detailed mechanistic theories of conscious access. The global neuronal workspace theory is usually presented by its proponents as a theory of access (Mashour et al. 2020). The hypothesis that it is

also a theory of phenomenal consciousness is a further step beyond the theory's core commitments (Carruthers 2019). More generally, the phenomenal/access distinction makes it possible to debate how phenomenality relates to access and, in particular, whether phenomenality "overflows" access. This has led to a rich and complex empirical controversy (Phillips 2018).

The conceptual split has, perhaps unintentionally, induced a subtle division of labour. Scientific work exploring the mechanisms of conscious access can now be conducted, and published, with little or no discussion of whether the work makes any progress towards understanding phenomenal consciousness. That issue can be left for articles reflecting on the wider implications of the scientific work. The adoption of the "conscious access" terminology by researchers in the global workspace program suggests they, at least, find this division of labour productive. Others who disavow that terminology (such as proponents of the integrated information theory) disagree.

Why might one find the division of labour counterproductive? Suppose one thinks that phenomenal consciousness should be the primary object of study of consciousness science, but suspects that access/availability of information is unlikely to be the basis of phenomenal consciousness. From this perspective, the construction of an "access consciousness/conscious access" concept is bad engineering: it has led scientific effort to be misdirected towards access/availability and away from the field's proper object of study. Moreover, it has encouraged a widespread presumption that the function of phenomenal consciousness must have something to do with access, a presumption that is only occasionally challenged (Scott et al. 2018). Worse, it has encouraged the view that phenomenal consciousness, being allegedly "distinct from any cognitive, intentional, or functional property" (Block 1995, 230), is beyond the reach of scientific inquiry altogether. That was not Block's intended message, but it is the message some have taken from his distinction. More speculatively, it may also have encouraged consciousness scientists to neglect the ethical implications of their work, a problem documented by Mazor et al. (2021). For it allows the thought: "phenomenal consciousness is what matters for ethics, but I study conscious access, so my work has no direct ethical implications."

Let us turn back to animal welfare. Human consciousness science gives us a case-study of how conceptual splitting can affect the trajectory of a research program, providing a basis for thinking about the risks and benefits of a split between "functional" and "experienced" welfare. On the plus side, we could expect to see progress in the direction of more detailed theories of functional welfare, and better-validated indicators of it. We could expect a division of labour to emerge between a body of empirical work measuring functional welfare, and a separate body of work about how functional welfare relates to experienced welfare (to some extent regimenting a division of labour that exists informally now). We could expect those immersed in research into functional welfare to find the division of labour productive.

On the negative side, if one thinks experienced welfare is the proper object of study of animal welfare, and that functional welfare matters only in so far as it correlates with experienced welfare (the view expressed by Duncan 1993), then one will naturally fear that the proposed conceptual split risks dragging the field away from its proper object of study. It will risk creating a default presumption that functional welfare tracks experienced welfare, discouraging serious empirical interrogation of this assumption. Worse, it will encourage the view that experienced welfare, being

allegedly distinct from any functional property, is beyond the reach of scientific inquiry. And it risks creating too much distance between the bulk of empirical research in animal welfare science and its normative mandate. If the field's central empirical project is guided by a welfare concept that is normatively inapt, that is still a serious problem, even if we retain a normatively apt concept for use in explicitly ethical and policy-relevant discussions.

Time to put my own cards on the table. I see the phenomenal/access split as a cautionary tale: on balance, I think the negatives have outweighed the positives. Human consciousness science should have accepted strain on its central concept as a sign of its immaturity as a scientific discipline rather than bifurcating the concept. Splitting the concept is a shortcut to easing strain, but a costly one. It is better on balance to retain a single concept and tolerate disagreement over its meaning than to split the concept in a way that creates a potentially counterproductive division of labour.

This is, I suggest, a lesson for animal welfare science, another young scientific field in the process of maturing. My tentative recommendation is that the field would be better served by retaining a single concept of animal welfare, defined partly in terms of consciousness—along the lines of Broom's concept—while recognizing that the demands of empirical validation and normative relevance currently put that concept under severe strain. Over the long run, we need to try to ease the strain by improving our empirical methods, so that the challenge from unconscious affect is no longer so severe. Splitting the concept is a tempting shortcut solution, but one that comes with significant risks.

Acknowledgements

I thank Kristin Andrews, Heather Browning, Andrew Crump, Simon Fitzpatrick, Simona Ginsburg, Eva Jablonka, Michael Mendl, Elizabeth Paul, Eva Read, Henry Shevlin, Walter Veit and audiences at PSA 2021 and the University of Edinburgh for discussion. This research is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, Grant Number 851145.

References

- Alexandrova, Anna. 2012. "Well-Being as an Object of Science." *Philosophy of Science* 79:678–89.
- Anders, Silke, Niels Birbaumer, Bettina Sadowski, Michael Erb, Irina Mader, Wolfgang Grodd and Martin Lotze. 2004. Parietal somatosensory association cortex mediates affective blindsight. *Nature Neuroscience* 7:339-340. <https://doi.org/10.1038/nn1213>
- Birch, Jonathan, Charlotte Burn, Alexandra Schnell, Heather Browning and Andrew Crump. 2021. *Review of the Evidence of Sentience in Cephalopod Molluscs and Decapod Crustaceans* (Report AW0517). Department for Environment, Food and Rural Affairs, UK. <https://www.lse.ac.uk/business/consulting/reports/review-of-the-evidence-of-sentiences-in-cephalopod-molluscs-and-decapod-crustaceans>
- Block, Ned J. 1995. "On a Confusion About the Function of Consciousness." *Behavioral and Brain Sciences*, 18:227-247.
- Boly, Melanie, Marcello Massimini, Naotsugu Tsuchiya, Bradley R. Postle, Christof Koch and Giulio Tononi. 2017. "Are the Neural Correlates of Consciousness in the Front or in the Back

- of the Cerebral Cortex? Clinical and Neuroimaging Evidence.” *Journal of Neuroscience* 37:9603-9613. <https://doi.org/10.1523/JNEUROSCI.3218-16.2017>
- Bornemann, Boris, Piotr Winkielman and Elke van der Meer. 2012. “Can you feel what you do not see? Using internal feedback to detect briefly presented emotional stimuli.” *International Journal of Psychophysiology* 85:116–124. <https://doi.org/10.1016/j.ijpsycho.2011.04.007>
- Broom, Donald M. 2008. “Welfare Assessment and Relevant Ethical Decisions: Key Concepts.” *Annual Review of Biomedical Sciences* 10:T79-T90.
- Carruthers, Peter. 2019. *Human and Animal Minds: The Consciousness Questions Laid to Rest*. Oxford: Oxford University Press.
- Chalmers, David. 1995. Facing up to the problem of consciousness. *Journal of Consciousness Studies* 2:200-19.
- Chalmers, David J. 2020. What is conceptual engineering and what should it be? *Inquiry*. Advance online publication, <https://doi.org/10.1080/0020174X.2020.1817141>
- Dawkins, Marian Stamp. 1988. “Behavioural Deprivation: A Central Problem in Animal Welfare.” *Applied Animal Behaviour Science* 20:209-225.
- Dawkins, Marian Stamp. 2012. *Why Animals Matter: Animal Consciousness, Animal Welfare, and Human Well-Being*. Oxford: Oxford University Press.
- Dawkins, Marian Stamp. 2021. *The Science of Animal Welfare: Understanding What Animals Want*. Oxford: Oxford University Press.
- Duncan, Ian J. H. 1993. Welfare is to do with what animals feel. *Journal of Agricultural and Environmental Ethics* 6 (supplement 2):8-14.
- Fraser, David. 2008. *Understanding Animal Welfare: The Science in its Cultural Context*. Chichester: Wiley-Blackwell.
- Korsgaard, Christine M. 2018. *Fellow Creatures: Our Obligations to the Other Animals*. Oxford: Oxford University Press.
- LeDoux, Joseph E. and Daniel S. Pine. 2016. “Using Neuroscience to Help Understand Fear and Anxiety: A Two-System Framework.” *American Journal of Psychiatry* 173:1083-1093. <https://doi.org/10.1176/appi.ajp.2016.16030353>
- LeDoux, Joseph E. and Richard Brown. 2017. “A Higher-Order Theory of Emotional Consciousness.” *Proceedings of the National Academy of Sciences USA* 114:E2016-E2025. <https://doi.org/10.1073/pnas.1619316114>
- Mashour, George A., Pieter Roelfsema, Jean-Pierre Changeux and Stanislas Dehaene. 2020. “Conscious Processing and the Global Neuronal Workspace Hypothesis.” *Neuron* 105:776-798. <https://doi.org/10.1016/j.neuron.2020.01.026>
- Mazor, Matan and 13 others. (2021). The scientific study of consciousness cannot, and should not, be morally neutral. *PsyArXiv*. <https://doi.org/10.31234/osf.io/8nb2c>
- Mendl, Michael, Oliver H. P. Burman, Richard M. A. Parker and Elizabeth S. Paul. 2009. Cognitive bias as an indicator of animal emotion and welfare: Emerging evidence and underlying mechanisms. *Applied Animal Behaviour Science* 118:161–181.
- Michel, Matthias. 2019. Fish and microchips: On fish pain and multiple realization. *Philosophical Studies* 25:95-110. <https://doi.org/10.1007/s11098-018-1133-4>
- Oizumi, Masafumi, Larissa Albantakis and Giulio Tononi. 2014. “From the Phenomenology to the Mechanisms of Consciousness: Integrated Information Theory 3.0.” *PLoS Computational Biology* 10: e1003588. <https://doi.org/10.1371/journal.pcbi.1003588>

- Paul, Elizabeth S., Shlomi Sherb, Marco Tamietto, Piotr Winkielman and Michael T. Mendl. 2020. Towards a Comparative Science of Emotion: Affect and Consciousness in Humans and Animals.” *Neuroscience and Biobehavioral Reviews* 108:749-770.
- Peters, Megan A. K., Robert W. Kentridge, Ian Phillips and Ned Block. 2017. “Does Unconscious Perception Really Exist? Continuing the ASSC20 Debate.” *Neuroscience of Consciousness* 2017:nix015. <https://doi.org/10.1093/nc/nix015>
- Phillips, Ian. 2018. The methodological puzzle of phenomenal consciousness. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 373: 20170347. <http://doi.org/10.1098/rstb.2017.0347>
- Salte, Liora. 1988. *Mandated Science: Science and Scientists in the Making of Standards*. Dordrecht: Springer.
- Schwitzgebel, Eric. 2016. Phenomenal consciousness, defined and defended as innocently as I can manage. *Journal of Consciousness Studies* 23: 224-235.
- Scott, Ryan B., Jason Samaha, Ron Chrisley and Zoltan Dienes. 2018. Prevailing theories of consciousness are challenged by novel cross-modal associations acquired between subliminal stimuli. *Cognition* 175: 169–185. <https://doi.org/10.1016/j.cognition.2018.02.008>
- Tamietto, Marco and Beatrice de Gelder. 2010. “Neural Bases of the Non-Conscious Perception of Emotional Signals.” *Nature Reviews Neuroscience* 11:697-709. <https://doi.org/10.1038/nrn2889>