



Is absence of evidence of pain ever evidence of absence?

Deborah J. Brown¹ · Brian Key²

Received: 18 June 2020 / Accepted: 16 November 2020
© The Author(s) 2021

Abstract

Absence of evidence arguments are indispensable to comparative neurobiology. The absence in a given species of a homologous neural architecture strongly correlated with a type of conscious experience in humans should be able to be taken as a *prima facie* reason for concluding that the species in question does not have the capacity for that conscious experience. Absence of evidence reasoning is, however, widely disparaged for being both logically illicit and unscientific. This paper argues that these concerns are unwarranted. There is no logical barrier to formulating cogent absence of evidence arguments; indeed, accepting such arguments is part of what it is to be committed to falsifiability as a critical aspect of the scientific method. Absence of evidence arguments can always be blocked, however, by assuming that psychological properties are ‘multiply realizable’. While we take multiple realizability to be highly likely at some level of analysis, we argue that it is question-begging to assume that it exists at every level of analysis, and thus it should not automatically be thought to undermine absence of evidence reasoning in the animal consciousness debate. Using the example of pain and focusing on homologies at the level of information processing tasks, we show how, in the science of consciousness, an absence of evidence might well serve as evidence of absence.

Keywords Pain · Nociception · Multiple realizability · Absence of evidence arguments · Neuroscience · Falsifiability

Granting Agency: Australian Research Council Discovery Project (DP200102909).

✉ Deborah J. Brown
deborah.brown@uq.edu.au

✉ Brian Key
brian.key@uq.edu.au

¹ School of Historical and Philosophical Inquiry, University of Queensland, Brisbane 4072, Australia

² School of Biomedical Sciences, University of Queensland, Brisbane 4072, Australia

1 Introduction

You are as certain as you can be that the stone paperweight on your desk does not feel pain even though you have no evidence of this. Indeed, the fact that there is no evidence is precisely the reason why you are so certain. Each of us is certain of our own capacity for pain and pretty nearly certain of the capacity of other human beings to feel pain. On the spectrum between stones and human beings lie creatures about which it is harder to judge with near certainty but about which we hope at least to be able to make informed judgements. As we approach single-cell organisms, we will inevitably have to rely on absence of evidence reasoning in ruling against a species' capacity for pain. Null hypotheses might seem dull in comparison to hypotheses claiming that protozoa or plants or, if one is inclined to panpsychism, every particle of matter, is conscious, but they play an important scientific role. In inferential statistics, an absence of evidence provides evidentiary support for the null hypothesis. Sometimes, and the science of consciousness is a case in point, an absence of evidence might well be evidence of absence.

In what follows, we focus our attention on pain not only for illustrative purposes but because pain is robustly coupled to emotion, motivation and behaviour in humans, and thus, in comparative studies, invites arguments from analogy based on multiple criteria. With respect to pain, there seems to be a standard argumentative approach that has been adopted for either accepting or rejecting the pervasive existence of pain in phylogenetically-diverse organisms. On the negative side, it is argued first that motor behaviours do not necessarily indicate feeling states such as pain, and second, that defined neural circuitry is necessary for such states, and in its absence, there can be no pain. These ideas, which we endorse, are closely linked to the well-established evolutionary and bioengineering principle that *structure-determines-function*. On the positive side, it is counterargued that the absence of evidence for a certain kind of neural circuitry is not evidence of an absence of pain. It exists and is somehow—although unknown—multiply realized in different species. The dual concepts of “absence of evidence” and “multiple realizability” have consequently become closely intertwined in these arguments. In this paper, we focus on both of these issues and their intersection with the aim of moving the debate over the pervasiveness of pain in animals closer towards some form of closure.

Pain research is notorious for making absence of evidence reasoning based on specific neurological criteria appear less conclusive. Point out that a species does not have a neocortex (necessary for pain in humans¹) and researchers will often reject the implication and look instead for some other criterion—neural or behavioural—as indicative of the species' capacity for pain. Different argumentative strategies are deployed at that point—some leaning toward the conclusion that

¹ The distinction between instinctive behaviours generated by the brainstem and conscious experiences generated by the cortex is well-documented (LeDoux 2012, 2013, 2014a, b; Rose et al. 2014; Scott 2014). Putative evidence from lesion studies of non-cortical pain (e.g. Damasio et al. 2012) is inconclusive, as Damasio and his colleagues admit, since it is based on cases of patients with only partial damage to critical pain sites, the insula, right and left SII and MCC (Denis et al. 2016; Lenoir et al. 2018; Garcia-Larrea and Mauguière 2018; Mazzola et al. 2019).

pain is “ambiguous” (Klein 2008); others toward the conclusion that pain has a different realizer in different species (“realizer functionalism”; McLaughlin 2006; Armstrong 1968; Lewis 1970, 1972; Kim 1998, 2005); and others towards the conclusion that pain is a higher-order functional property—the property of having some first-order realizer that fulfils the functional role of pain (“role functionalism”; McLaughlin (2006)). What motivates such moves is an interwoven set of assumptions: that all genuinely causal or explanatory properties are functional (Shoemaker 1984, 1998, 2001); that statements about unobservable psychological events need to be “Ramseyfied” to be rendered scientific—i.e., translated into sentences existentially quantifying only over observable properties (typically observable causes and effects of the psychological event) and individuals (Lewis 1970, 1972); and the idea that, for any function, there is an open disjunction of actual or possible realizers of that function—the “multiple realizability” assumption (Putnam 1967). Viewed through the lens of comparative neurobiology, these assumptions, if true, rule out an absence of evidence hardly ever lending support to the null hypothesis about a species’ capacity for pain so long as there is some observable mapping of inputs (noxious stimuli) to outputs (nocifensive behaviour). This, we shall endeavour to show, is deeply problematic.

In Sect. 1, we defend absence of evidence reasoning drawing on, among other things, results from informal logic. These results delineate the kinds of epistemic contexts in which absence of evidence arguments can be cogent. This gives us a *prima facie* reason for thinking that we should be able to utilize absence of evidence arguments in the animal consciousness debate. In Sect. 2, we show how these epistemic conditions could be met in practice, drawing on extant investigations in comparative neurobiology. In Sect. 3, we draw out the implications of this investigation for the widely held assumption that psychological states like pain are multiply realizable. We thus seek to add to the voices of dissent advocating a revision of the multiple realizability assumption. In Sect. 4, we develop the idea of a common neural architecture across extant species that could serve as a minimal condition for a certain kind of conscious experience. We advocate theorizing about multiple realizability in observance of the foundational axiom of biology that *structure-determines-function*. Paying close attention to this axiom forces us to revise whether ‘realization’ is the right way to characterise the relationship between entities at different levels of analysis. Overall, if we are right, an absence of evidence of a certain kind of neural architecture could well be evidence of an absence of consciousness.

Let us be clear about our intentions: We do not assume that the example we offer below of an absence of evidence argument—‘FISH’—is sound; only that it is a kind of argument that *could* be cogent if the right epistemic conditions were met. We defend only the right to use the form of argument instantiated in FISH when reasoning about a given species’ capacity for conscious experience. Whether arguments of this form yield conclusions that are likely to be true is an empirical matter, and while we cite here some evidence for (and against!) FISH, this is only to provide a sense of what it would be like to go through the motions of evaluating an argument like FISH on the basis of scientific evidence (or a lack of it), rather than simply dismissing the argument for ideological reasons. The point of the exercise is to establish the acceptability of the form of the argument, not, on this occasion, its content.

2 Absence of evidence arguments and the science of consciousness

FISH: Animals that feel pain possess neural circuitry enabling them to execute the neural computations that lead to pain. There is no evidence that fish possess such circuitry. Hence, fish do not feel pain.

FISH is an instance of a form of argument that appears in neuroscientific debates about the capacity for conscious experience across phylogenetically-diverse species. In practice, the neuroscientist purports to have discovered a kind of circuitry that they take to be strongly correlated with or explanatory of a certain kind of conscious experience in humans, positively verified through first-person reports, and then reasons inductively to other species' capacity for that experience based on the presence or absence of homologies of that circuitry. This is standard evolutionary biological practice (Bock and Von Wahlert 1965; Wake 1992), and yet while many are inclined to accept positive evidence as indicative of the capacity for consciousness, few are inclined to accept negative findings as evidence of a lack of consciousness. FISH is an 'absence of evidence' argument. It reasons from negative findings to the conclusion that fish are incapable of feeling pain. Why should there be such asymmetry in the way positive and negative findings are regarded?

Among some researchers, the battle cry 'Absence of evidence is not evidence of absence!' has taken on the status of an axiomatic truth. Brown (2016b, p.2) writes: "It is true that "human-like" nociceptors have not been identified in elasmobranchs yet, but that is not to say that they do not exist. Science 101: Absence of evidence is not evidence of absence." In an article entitled "Assessing negative and positive evidence for animal pain", Elwood (2017) similarly objects to anyone's relying on negative evidence to conclude that a species is incapable of pain. While he may be correct that negative evidence resulting from an insufficiently extensive search is not reliable, his own reliance on the slogan "absence of evidence is not evidence of absence" (2017, p. 1) goes beyond this otherwise minor methodological cautionary tale. Manzotti (2018, p. 2) offers a stronger blanket prohibition of the use of absence of evidence arguments, stating that "there is no evidence that an animal that shows pain-avoidance does not feel pain, no matter what neural structure is involved", meaning by "no evidence" here an absence of evidence and assuming that no such negative findings could ever be conclusive of an absence of pain. Such aversion to absence of evidence reasoning is not restricted, moreover, to individual researchers. A report of the National Research Council of Canada, (2009, p. 23) states that:

Although definitive evidence is often unavailable, this report does not treat the absence of evidence as evidence of absence. Instead, the consensus of the committee is that all vertebrates should be considered capable of experiencing pain.²

Our task here is try to move this debate beyond its uncritical reliance on this slogan. Grounds for assuming that the slogan is true are rarely cited, but we can make some

² See <https://www.ncbi.nlm.nih.gov/books/NBK32655/#ch1.s5>.

educated guesses. Some might fear that absence of evidence reasoning is logically illicit. Those operating within a Bayesian framework might, for example, find it questionable to update beliefs on the basis of an absence of evidence. Some have thought that an absence of evidence is at best “weak” (Sober 2009) or “subjective” (Strevens 2009) evidence. It is true that an absence of evidence becomes weaker the more that *finding* evidence given the truth of a hypothesis is left to chance. If, for example, the probability of finding an intermediate fossil (Sober’s interest case) is low, then the absence of evidence will not warrant believing that intermediate fossils do not exist or in any way vindicate Creationism. Generally, however, there is no barrier within a probability framework to updating on the basis of an absence of evidence any more than on the basis of positive findings. An absence of evidence can form part of the ‘total evidence’ to be considered. It can tip the balance in favour of one hypothesis over a competing one, especially if the probability of finding evidence if a certain hypothesis is true is reasonably high and the evidence is not forthcoming. And an absence of evidence itself could be something that stands in need of a scientific explanation. Indeed, it could be the very thing that calls for a replacement or revision of an accepted theory or hypothesis. An absence of evidence should not, in other words, be thought of as necessarily epistemically inert.

FISH might appear logically spurious for other reasons, however. Superficially, it looks like an *argumentum ad ignorantiam* (argument from ignorance), which, in classical logic textbooks, would standardly be classified as a fallacy. Arguments from ignorance are bad because for any proposition, p , not knowing that p is true is not a reason for inferring that not- p is true (Walton 1996, p. 143). But FISH need not be so flimsy. If FISH were based on a complete lack of information about the neuroanatomy of fish, then it would be a bad argument. In that situation, negative results would at best provide a reason to suspend judgement or prompt further investigation, but never to draw a definitive conclusion in the way that FISH does. Not accepting the proposition that fish feel pain on the basis of an absence of evidence is not the same thing as accepting the proposition that fish do not feel pain, and, you might think, we should be careful to avoid conflating the two.

The central premise of Douglas Walton’s (1996) book, *Arguments from Ignorance*, is that arguments from ignorance are *presumptive* arguments—arguments that may be fallacious in certain epistemic contexts, but cogent in others. Consider the following two acceptable reasoning scenarios:

KEYS: You are looking for your keys and think you might have left them on the bookshelf, but when you look, you don’t see them. A natural conclusion to draw is that the keys are not there.

VULCAN: You’re an early twentieth century astrophysicist seeking to test Le Verrier’s hypothesis that there is a planet (Vulcan) causing perturbations in Mercury’s orbit. You keep looking but find nothing. You conclude that Vulcan does not exist. (If you happen to be Albert Einstein, you devise an alternative explanation—General Relativity theory—live long, and prosper.)

These are absence of evidence arguments, and yet are cogent.³

Ultimately, what all these considerations suggest is that absence of evidence arguments are neither cogent nor fallacious in themselves but only in relation to their contexts of use. In KEYS and VULCAN, what makes the arguments cogent is that in each context the agent could reasonably have expected to find evidence if the hypothesis were true, and yet did not. This is sometimes referred to as the *epistemic closure principle*. Put simply: ‘*If it were true, I would know it*’ (de Cornulier 1988, p. 182; Walton 1996, p. 147).⁴

Armed with the epistemic closure principle, we can see an obvious path to making FISH not just a rationally acceptable form of argument but one that is, even by the standards of classical logic, valid. Indeed, it would be difficult to dismiss FISH as illogical for the simple reason that it can be easily transformed into a logically valid form of argument, *modus tollens*:

Modus Tollens:

(P1) If p, then q.

(P2) Not-q

(C) Therefore, not-p.

Adding a tailored version of the epistemic closure principle (P1* below), we can transform FISH into FISH*, a modus tollens argument:

FISH*:

(P1*) If the hypothesis that fish feel pain were true, diligent investigation would identify neural circuitry enabling fish to execute the neural computations that lead to pain.

(P2*) Diligent investigation does not identify neural circuitry enabling fish to execute the neural computations that lead to pain.

(C*) Therefore, the hypothesis that fish feel pain is not true.

We have chosen the non-probabilistic form (FISH*) for ease of exposition. Because FISH* is a presumptive, scientific argument, however, it might be better to construe it probabilistically, as in FISH**:

³ While some absence of evidence arguments are cogent, there are clear cases of fallacious arguments that assume that an absence of evidence is not evidence of absence. At a NATO press conference in 2002, Donald Rumsfeld declared the invasion of Iraq justified on the grounds that although there was no evidence Iraq had weapons of mass destruction, “Simply because you do not have evidence that something exists does not mean that you have evidence that it doesn’t exist.” (<https://www.nato.int/docu/speech/2002/s020606g.htm>) More recently, physicists, Marcel Gleiser and Martin Rees have each won Templeton prizes for “exceptional contributions to affirming life’s spiritual dimension,” Gleiser, for putting the proverbial catechism among the dogmatists and declaring atheism unscientific on the grounds that an absence of evidence of God is not evidence of God’s absence, and Rees, by affirming the probability of undiscovered terrestrial super-intelligences and intelligent extraterrestrials on grounds including that an absence of evidence is not evidence of absence. There is perhaps at least as much danger in relying on the slogan as on its negation. See Brown and Key (2019) for discussion of these considerations.

⁴ To see how this works, consider Gleiser’s argument in the previous note. Most theists would reject Gleiser’s argument against atheism precisely because the assumption upon which it depends, namely, that God is *Deus absconditus* (a “hidden God”) is explicitly rejected by most religions. One exception is Lutheranism. Most religions proclaim that God reveals himself to his followers, from which it follows that if the hypothesis that God exists is true, the probability of finding evidence that God exists should be high. If Gleiser is right, and there is instead an absence of evidence, then atheism is vindicated, not refuted.

FISH**:

(P1) If the hypothesis that fish feel pain is true, then it is highly probable that diligent investigation would identify neural circuitry enabling fish to execute the neural computations that lead to pain.

(P2) We are (reasonably) certain that diligent investigation does not identify neural circuitry enabling fish to execute the neural computations that lead to pain.

(C) Therefore, it is highly probable that the hypothesis that fish feel pain is not true.⁵

As Sober (2009, p. 64) observes, one problem with adopting this kind of strategy to justify absence of evidence reasoning is that in some cases the epistemic closure principle is true; in other cases, not. Sober's counterexample involves the following scenario. Suppose it happens to be true that there are storms now on Jupiter (but we don't know it). The likelihood of finding evidence of storms on Jupiter right now is, however, low. In this case, we could not reasonably expect to find evidence right now even if the hypothesis were true, and so the fact that we do not have evidence in support of the hypothesis is no reason to reject it. (Of course, it is no reason to accept it either!)

This tells us two things. As Sober (*ibid*) points out, an absence of evidence doesn't logically entail an absence—it is evidence of absence except when it isn't! But it also tells us something else about the epistemic conditions under which an absence of evidence is evidence of absence. Besides closure, we need sufficient "epistemic coverage" (Goldberg 2011; Pedersen and Kallestrup 2013). KEYS and VULCAN work because in both contexts our investigations are extensive and the contexts are restricted enough for us to draw a 'positive- negative' conclusion (i.e. that what we are looking for is not there). Epistemic coverage is sometimes construed in reliabilist terms—i.e., agents have good epistemic coverage of the evidence when they are connected to a source that feeds them reliable information about the domain and where they are receptive to being "hooked in" to information from that domain (Pedersen and Kallestrup 2013, p. 2576). But whether one needs to meet reliabilist standards in order to claim epistemic coverage is a matter which need not concern us here. The important point for our purposes is that in order to judge the cogency of an absence of evidence argument, there must be some mechanism available for evaluating the degree of epistemic coverage. We cannot assume right now to have good epistemic coverage of the surface of Jupiter, but the neuroanatomy of fish is a different kettle of fish. Whether FISH* is sound or not would depend on whether the epistemic closure principle (P1*) and the epistemic coverage claim (P2*) are true. So far though, nothing about the form of the argument prevents it from being an argument we could rationally accept.

⁵ The cogency of probabilistic modus tollens is ably defended by Wagner (2004) and Sobel (2009).

3 Absence of evidence of neural circuitry in fish

What we learned from the previous section is that where absence of evidence arguments work, they do so on the basis of a mix of negative findings and positive knowledge (Copi 1982; p. 102; Walton 1996, pp. 135–138). The epistemic closure principle is a positive claim, which, when supported by a true premise asserting an absence of evidence, provides grounds for inferring that the hypothesis that something exists is false. What would it be like to have such grounds for accepting the conclusion of an absence of evidence argument in relation to the science of consciousness?

Let's go through the exercise of evaluating FISH (or its more formal variations, FISH* or FISH**). The first task would be to establish that the epistemic closure principle (P1*) in FISH* is true.

We start by defining pain as the conscious (felt) neural processing of noxious stimuli, in line with the standard neuroscientific definition (Anand and Craig 1996; Williams and Craig 2016) that emerged from the experimental work of C.S. Sherrington in the late 1800s (Woodworth and Sherrington 1904). This definition implies a distinction between pain and nociception, the non-conscious (unfelt) processing of noxious stimuli and associated nocifensive behaviours (e.g. the flexion withdrawal reflex). Given the overwhelming evidence in favour of the dissociation between pain and nociception,⁶ and their being seated in different neural circuitry (Sherrington 1947; Geuter et al. 2020; Tracey 2005), we would expect this distinction to be uncontroversial. But it is not, creating enormous confusion.

It is the lack of attention to this distinction that lends more credence than deserved to the alternative functionalist solutions outlined earlier. Functionalists typically allow that fish feel pain even though they lack the neural circuitry evident in humans and other mammals in one of three ways: (1) fish have their own kind of pain and 'pain' is ambiguous (Huntingford et al. 2006; Sneddon 2015; Seth 2016; Manzotti 2018); (2) pain in fish is the same kind of pain as that experienced by humans but has a different neural realizer (Lewis 1972, 1980); or (3) 'pain' refers to a relational, higher-order functional role property shared by fish and humans but executed by different mechanisms in each species. None of these options does much to preserve a robust distinction between pain and nociception. (1) makes it hard to see why any kind of nocifensive behaviour, including that exhibited by single-celled organisms, would not count as a 'sort of pain'.⁷ (2), the "Ramsey-Lewis" solution, prioritizes behavioural criteria resulting in a kind of 'black box functionalism' that renders the approach insufficient to dismiss the null hypothesis. There is nothing, for example, in Lewis' story of Martian pain, based as it is on observations of behaviour, that favours the hypothesis that the Martian is in pain over the null hypothesis that the Martian merely has nociception. (3) assumes that the functional

⁶ In humans, there may be either pain without nociception and nocifensive behaviour (such as in central pain syndromes; Loeser and Melzack 1999) or nociception with behaviour and no pain (Dimitrijevic and Nathan 1968). There are also cases of the behaviour normally associated with pain—as in pathological crying in pseudobulbar affect (Parvizi et al. 2006)—with neither nociception nor pain.

⁷ See calls for recognizing pain as an evolutionarily required feature of all living creatures (Baluška and Reber 2019).

role is itself causally explanatory, leaving it open to the objection that what is really doing the causal work is the particular neural realizer (Block 1990, p. 155; Kim 1998, p. 51). Since (3) places no constraints on what kind of thing is required to perform the specific function of pain, it too is insufficient to maintain the distinction between pain and nociception.

The costs of conflating pain and nociception are high. Consider the first option (1)—that ‘pain’ is ambiguous. Let’s call the kind of pain fish have on this view ‘pain \neq ’ to distinguish it from human pain. The problem is that we do not know what pain \neq is like and so do not know what, beyond crude similarities in behaviour, to look for as the structure that explains pain \neq . We could not assume that *what-it-is-like* (Nagel 1974) for fish to feel pain \neq , the “qualia” or “phenomenal” aspect (Block 1980; Jackson 1982) of their conscious experience, would be the same as pain experienced by humans. If we could, ‘pain’ wouldn’t be ambiguous. Strictly speaking, it is a category mistake to say that pain \neq is painful, since the predicate ‘is painful’ inherits its meaning from ‘pain’ not ‘pain \neq ’. But now it is not clear what we are talking about when we refer to ‘fish pain \neq ’ or how we could establish that pain \neq is even a feeling rather than just another name for nociception. If the criteria for both are the same, any pain \neq hypothesis would be unverifiable. We could instead regard pain and pain \neq as sharing the same epiphenomenal, emergent property, which, as Kim (1998 p. 18) notes, alleviates the problem by allowing type-identical epiphenomenal properties across different realizers with different causal powers. But this seems to us to make the question of realization moot, since what makes each a kind of pain is the shared epiphenomenal property and not what physically causes or constitutes pain. This has to be the counsel of despair.

(2) and (3) do not fair much better. The problem with the Ramseyfication in (2) or focusing on what the psychological state does (observationally) at the expense of what it is, as in (3), is that we lose sight of the intermediate, computational level of analysis. This level is couched in irreducibly theoretical language and involves abstracting from descriptions of lower-level constitutive properties. Without access to this level, it will always be tempting to obfuscate the distinction between pain and nociception, rendering any hypothesis that a creature has pain unverifiable as the criteria for being in pain become indistinguishable from the criteria for exhibiting nociception.

An alternative to all three options outlined above is to suppose that ‘pain’ is unambiguous and that it indeed refers to a “nomically projectable kind” (Kim 1998, p. 109). On this approach, we would seek to identify the neural architecture that explains why an organism has the phenomenal experience it does and whatever functional role that experience typically (but perhaps not universally) performs. Such an approach would yield conditionals of the form ‘If x is in pain, then x has neural circuitry, N ’. The consequent would describe a *necessary* condition for pain that could serve to ground predictions of different species’ capacity for pain. If fish and humans both experience pain (*i.e.* both feel the same negative affect valence), then we could expect that discoveries about how pain is manifested in humans would inform decisions about whether fish feel pain. If we found similar circuitry, we could then explain why both fish and humans writhe about when administered electric shocks and how they use their pain to modify their behaviour so as to avoid

noxious stimuli. However this investigation pans out, we could assume all the while that pain constitutes a single, unified, scientific kind while preserving the distinction between pain and nociception, on the one hand, and avoiding epiphenomenalism, on the other. Happy land!

This outcome would find considerable support within the community of neuroscientists claiming to have found evidence of nociception in fish but no pain (e.g., Adamo 2018; Derbyshire 2016; Favela 2017; Gamez 2018; Gerlai 2017; Gutfreund 2017; Rose et al. 2014; Stauffer 2017; Key 2015, 2016; Key and Brown 2018). It is well known that pain in humans is dependent on neural activity in the cerebral cortex of the dorsal telencephalon (Key and Brown 2018). On this basis, we could reasonably expect to find structurally homologous brain regions in fish that are also necessary for pain. However, surgical removal of the whole dorsal telencephalon (called the ‘pallium’) in fish fails to abolish the avoidance behavioural responses to noxious stimuli (in this case, electric shock; Portavella et al. 2004). Whatever is happening in the dorsal telencephalon of fish then is not responsible for its nocifensive behaviour and thus does not support the presence of pain. The absence of any signature pain-relevant neural architecture in fish would lend credence to the second premise of FISH*, P2*. If the neuroanatomy of fish were well-mapped, (as we have argued to be the case; Key 2015, 2016), indicating good epistemic coverage of the domain, by *modus tollens*, the conclusion that fish lack the capacity to feel pain would go through.

This is what it would be like to conclude that a species is not capable of pain on the basis of an absence of neurological evidence and plausible functionalist assumptions. But, of course, the elephant in the room is the prospect of multiple realizability, which we have conveniently so far ignored. The multiple realizability principle can always be trotted out to defeat an absence of evidence argument. Whether it should is another matter.

4 Addressing multiple realizability at the correct level of abstraction

It is important not to underestimate the role that the multiple realizability assumption plays in blocking absence of evidence arguments in the animal consciousness debate. When challenged to explain how a creature is able to feel pain when it lacks the neural architecture responsible for pain in humans, scientists and philosophers alike often defer to the concept of multiple realizability (Demski 2013; Sneddon 2015; Broom 2016; Brown 2016a; Merker 2016; Seth 2016; Striedter 2016; Woodruff 2017, 2018; Michel 2019). Pain, it is said, is a function and functions are essentially multiply realizable (Putnam 1967). Just as chairs involve many different kinds of structures and materials, so too it is argued that pain is realized in different kinds of neural circuitry, neural configurations, or the green goo of the inflated feet of Martians (Lewis 1980). There is no monolithic attachment to multiple realizability, however. Polger (2004, p.11) points out that we either have to reject multiple realizability or deny that we have psychological experiences in common with other species. He speculates that the identity theorist might indeed find commonalities between all critters who experience pain, noting, however, its

status as an “ambitious hypothesis” (Polger 2015, p. 874). When Haikonen (2016) asks whether we should not just assume that the neural realizer for pain is a “simple one” that could fit into either “the diminutive fish brain or a complex one calling for larger cortical resources?”, he is swallowing the “ambitious hypothesis” hook, line and sinker. No one has any idea what this “simple realizer” might be.

One source of support for the multiple realizability assumption are folk arguments from analogy—jet planes fly without propellers (Ng 2016), crickets hear without human ears (Dinets 2016), teeth and gastric acid perform the “same function” (Jones 2016), and bats fly without feathers (Manzotti 2016). Sameness of function is possible in each of these examples only because of the indeterminate way in which the function is defined. For the concept of flying, for example, it is possible to imagine many different objects ‘flying’, including tennis balls, arrows, rockets and birds. If, instead, we are more discerning and consider flying at a more discriminatory and explanatory level—such as one involving aerodynamics—the argument for multiple realizability soon breaks down. At this level, the forces of lift, gravity, thrust and drag are markedly different between disparate structures. Honing in on a more determinate function—flying based on the capacity for aerodynamic lift—structures begin to exhibit more commonalities than differences. Despite variation in the shape and form of wings or wing-like appendages, all species capable of performing this function exhibit a common wing design that explains their capacity for aerodynamic lift (Lindhe Norberg 2002). There are clear explanatory benefits of this more fine-grained approach to functional categorization as it allows us to see what it is about the structure of the wing that explains its capacity for aerodynamic lift. Of course, one can swing too far in the opposite direction—categorising the function of being consciously aware of a red light in the left side of the visual field as being distinct from that of being aware of the same light coming from the right. Choosing the appropriate level of analysis and description for both structure and function is thus critical to the evaluation of whether there is multiple realizability of a function or not. When either the structural level is too granular or the functional level too indiscriminate, multiple realizability is assured—and its argumentative value is exhausted.

To get the structure and the function, and indeed, the level of analysis, aligned is not easy. Citing numerous studies revealing common neural substrates for visual processing in the brains of humans and monkeys, Bechtel and Mundale (1999) reject multiple realizability. But, of course, within any brain region there may be much variation. Beyond establishing correlations between cognitive functions and brain regions, neuroscience endeavours to decipher what it is about the structures that distinct species have in common that explains their ability to perform the cognitive function under examination. This imposes a tighter constraint on the attribution of a cognitive function than establishing correlations to brain regions. Specifically, neuroscience looks beyond homologies of brain regions to homologies of information processing tasks. This requires attention to the neural architecture or precise circuitry that makes possible the performance of a function.

This insistence on paying attention to intermediate levels of analysis and structure is not new (Nelson 1969; Piccinini and Maley 2014; Polger and Shapiro 2016; Bickle 2003; Wimsatt 2007; Craver 2007). We agree with Piccinini and

Maley (2014, p. 131) that one needs to isolate the “relevant causal mechanisms” in judging whether there is multiple realization or not and with Polger and Shapiro (2016, p.98) that “sameness of function owes to sameness of structure.” In regard to neuroscience, however, a common problem is that theorists tend to operate at too high a level of generality. Polger and Shapiro, for example, never interrogate what kinds of circuitry underlie and explain the functions of memory, pain, trichromacy, etc., that they describe. In the end, they too fall back on behavioural criteria in making the call about species with neural circuitry different from that presented by humans—e.g. when considering whether octopi feel pain, they write that “the right conclusion is that...their pain experiences differ from human experiences” (2016, p. 113; cf: Key and Brown 2018). The result is some unfortunate mixed messaging. On the other side, Gillett (2002; 2007, 2010) and Aizawa and Gillett (2009), defend a “dimensioned” framework for attributions of multiple realizability that attempts to move away from reliance on intuitions. Multiple realizability applies only in contexts of hierarchically and compositionally organised entities, their properties (especially, “causal powers”) and relations. To avoid triviality, examples of multiple realizability must be drawn from the same level of analysis. But in the end, multiple realizability is guaranteed since one will always find some differences among the realizers. As Piccinini and Maley (2014, p. 128) point out, not every difference in the realizer interferes with its capacity to realize a given property, or hence, counts as a case of multiple realization.⁸ In neuroscience, different protein structures and different sizes and numbers of subcellular compartments like dendritic spines that in turn create fluctuations in the excitatory electrical signalling of neurons, show how neurons multiply realize “many higher-level psychological properties” (Aizawa and Gillett 2009, pp. 200–201). This, we agree, is multiple realizability; it is just that it does not exclude non-multiple realizability at a level which is more explanatorily relevant.

Assessing multiple realizability in neuroscience requires paying attention to the correct level of abstraction. Structural differences at the levels of gene expression, subcellular compartments (e.g. synapse density, dendritic tree branching) and neuron population number are too granular, whereas the mere presence of a forebrain is too gross for drawing meaningful conclusions. Piccinini and Maley (2014, p. 144) touch on this idea with respect to action potentials in the nervous system, acknowledging that multiple realizability occurs only at some levels of organisation. They claim that a neuron’s capacity to fire an action potential is not multiply realized even though the ion channels producing the action potential could be. That is, action potentials could only be realised by the flow of ions across the plasma membrane, and while there may be various kinds of ion channels with

⁸ Piccinini and Maley (2014, p. 135) regard the realized properties of a whole “considered in abstraction from its organised parts” as “proper subsets of the causal powers of their lower-level realizers” (ibid, p. 147–148), and thus while not autonomous from their lower-level realizers, are not reducible to them either. We wonder how the relevant subsets of causal properties are identified, preferring instead to work with the notion of a *structure*. In the case of neuroscientific explanation, the relevant structure is the circuitry that explains the possibility of the organism performing the neural computations necessary for a given type of conscious experience. We also prefer to rely on different types of inter-level relationships besides realization, as we discuss below in Sect. 4.

relevantly different molecular properties, each of these kinds could still contribute to generating ion flow. We highlight this problem by examining a simple neural circuit—the central pattern generator (CPG)—that drives rhythmic locomotion in bilaterally symmetrical animals. The function of the CPG is to generate rhythmical bursts of action potentials that are causally responsible for cyclical movements of either body segments or limbs on either side of the body, as in stepping or swimming. The structure and function of CPGs have been well studied in Mollusca and, in particular, in two closely related nudibranchs—*Melibe leonina* and *Dendronotus iris*. The gross swim behaviour of these two species appear very similar and is described as stereotyped rhythmic left–right movements consisting of simple C-shaped or lateral bends of the body occurring at a frequency of three cycles per second (Sakurai et al. 2011). If we assume that this function is identical between species, nudibranchs provide a convenient model for assessing whether the CPGs have similar neural architectures, especially since these closely related species have many anatomically homologous neurons.

The CPGs in these two species share similarities but are not identical (Sakurai and Katz 2019; Fig. 1). In *Dendronotus*, there are two swim interneurons—Si2 and Si3—that are present on each side of the body that constitute the CPG. In contrast, *Melibe* has the same two homologous neurons as well as an additional two pairs—Si1 and Si4—that form the CPG. In addition to the extra swim interneurons and their interconnectivity, the synaptic connectivity between the common neurons—Si2 and Si3—is disparate.

Neither neuron number, synapse number nor specific axonal interconnectivity appear to be suitable structural parameters for characterizing CPG between these species. However, on closer inspection, it is apparent that both species share the common structural feature of reciprocal crossed inhibition of interneurons. The projectable kind here is determined not by the number of individual neurons, which varies (and so there is multiple realizability at this level), but by the simplified structure, which describes the minimal conditions for left–right locomotion. Reciprocal inhibition has long been considered to be essential for the normal execution of locomotor behaviours in bilaterally symmetrical animals. This can be demonstrated by experimentally inducing unilateral blockage of crossed inhibitory neurons (Moult et al. 2013). When ipsilateral inhibitory neurons are prevented from inhibiting their contralateral counterparts, there is loss of rhythmic locomotor activity. Thus, both species possess neurons that project to the opposite side of the body to inhibit the contralateral neurons in the CPG. This neural connectivity is essential to ensure that one side of the body is relaxed while the other side contracts to perform the rhythmical locomotion—in this way, the structure of the circuitry explains its functional role. To date, all species exhibiting rhythmical locomotion employ crossed inhibition using crossed axons that lead to contralateral inhibition of motoneurons in the context of the CPG (Katz 2016; Kiehn 2016; Harris-Warrick and Ramirez 2017). The presence of neurons performing this neurocomputation becomes a useful criterion for determining whether an animal has the potential for performing rhythmical locomotion. At the level of analysis involving the necessary neurocomputation of crossed inhibition—the *explanatory level*—multiple realizability is no longer applicable. This strongly suggests that multiple realizability does

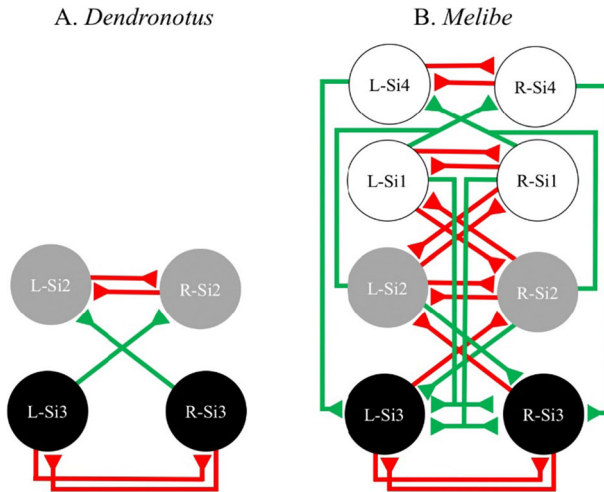


Fig. 1 Simplified circuitry diagrams of central pattern generators in **a** *Dendronotus* and **b** *Melibe*. Circles represent neurons which are referred to as swim interneurons (Si) 1–4 on either the left (L) or right (R) sides of the midline. Red and green lines represent axons with either inhibitory or excitatory monosynaptic chemical synapses, respectively. (Color figure online)

not automatically apply at all levels of analysis, and therefore must not be assumed uncritically in rejecting absence of evidence arguments like FISH*.

While we have argued here that choosing the correct level of abstraction is critical for determining whether a function is multiply realized, we have not, as yet, provided any clue as to how one goes about identifying this level—apart from suggesting that the analysis should be neither too granular nor too obtuse.⁹ Are there any rules or heuristics that could inform this choice? In this matter, we can be guided by the earlier work of Marr (1982) on the visual system. Marr reasoned that to understand the function of a nervous system, one has to address three levels of analysis: Level 1—the computational level that describes the goal of the system; Level 2—the algorithmic level that details the process of performing specific computations to achieve the system’s goal; and Level 3—the implementation level that contains the basic components that execute the algorithm. Marr realized that a computational goal could possibly be achieved via different algorithms in the same system and that the same algorithm may be physically realised by different circuits in different systems. Thus, the algorithmic level is probably not the best initial choice for a level of abstraction. By ‘computational goal’, Marr was referring not to the overall goal of visual perception, but rather to specific computational tasks such as stereopsis and motion detection that are necessary for normal human vision. Perturbations to relevant brain regions executing motion detection lead to perturbed visual perception (Zeki and Bartels 1998). Visual motion detection is also an essential computational goal in many visual systems and is realized by discrete neural modules in widely diverse species such as humans, insects and crustaceans

⁹ Disambiguating the notion of ‘levels’ is a further matter for contention. See Wimsatt (1994, 2007), Craver (2007) and Barberis (2017).

(Strausfeld 2005; Borst and Helmstaedter 2015; Strother et al. 2017). In the absence of these modules, motion detection and associated behaviours are compromised. High level computational goals (i.e. specific information processing tasks), therefore, appear to be an appropriate starting point for selecting the correct level of abstraction.

While computational goals can be elucidated by theoretical approaches, their necessity must subsequently be confirmed through experimental manipulation and/or clinical case reporting—as illustrated above for bilateral symmetrical locomotion and motion detection. A number of specific information processing tasks have been shown to be essential for conscious processing of noxious stimuli (Key 2015, 2016). This does not prescribe how conscious awareness is generated by these computational goals, but only that these processes form the essential building blocks for pain. For example, pain depends at least on parcellation of neural tissue into modules or domains that perform pain-dependent computations accounting for the intensity (Porro et al. 1998), quality (Maihöfner and Kaltenhäuser 2009), unpleasantness (Rainville et al. 1997), motivational value (Wager and Barrett 2017), and localisation (Baumgärtner et al. 2006) of a noxious stimulus. These regions have been shown to have strong reciprocal interactions involving inhibitory and excitatory interconnections to make this possible. Supposing that the presence of such discrete brain regions and their feedforward and feedback interconnections are essential for pain, it would be reasonable to investigate whether these neural regions and circuits were realised in the fish brain in any form (*i.e.* without being restricted by phylogenetic homologies of brain regions or without concern for low level structural features such as local microcircuitry, neuronal population size and subcellular features). There is evidence that most of the forebrain in fish is not necessary for behavioural responses to noxious stimuli, and that the sub-forebrain regions driving this behaviour (as determined by recent whole brain imaging and ablation studies in zebrafish; Lovett-Barron et al. 2019; Wee et al. 2019) lack the neural architectures considered necessary to execute the prerequisite information-processing tasks for pain (Key and Brown 2018).

By focusing on the necessity of specific information-processing tasks for perception rather than on multiple ways of organising processing steps within an algorithm, we can see how the debate could move beyond the simple assumption of multiple realization. While we may continue to debate whether a particular information-processing task is necessary or not for pain, this does not negate the idea that there are certain core computations that are essential for pain. It merely means that those neural computations become a matter for empirical verification. As recently demonstrated, this approach is providing important insights into the evolutionary origins of phenomenal consciousness (Key and Brown 2018).

5 Beyond multiple realizability

The aim in the previous section was to defend the legitimacy of absence of evidence arguments in comparative neurobiology by arguing against the multiple-realizability-at-every-level assumption. We recognise, however, that many theorists would die in a ditch before jettisoning multiple realizability. Some, (e.g., Allen 2013; Jones 2016), go so far as to claim that pain might be something better investigated at the “behavioural and cognitive levels”:

The possibility of conscious experience at the behavioural and cognitive levels despite morphological and anatomical differences at the neurological level makes fish an enormously interesting testing ground for ideas about multiple realizability of cognition. (Allen 2013, p. 36¹⁰; cited in Jones 2016)

This is a strange call to arms. It entails that in investigating whether a species feels pain, we should ignore state-of-the-art neural recording and imaging, molecular genetic manipulations allowing for the dissection and identification of distinct neural circuits, electrophysiological and advanced microscopy techniques involving neural circuit manipulation and mapping—just to maintain an attachment to multiple realizability. How might we better align our philosophical intuitions with such state-of-the-art advances in neuroscience?

One diagnosis for this misalignment is that theorists have become overly fixated on the realization relation. Because they understand the realizer relation compositionally, Aizawa and Gillett (2009, p. 194) insist that it is transitive. If *x* realizes *y* and *y* realizes *z*, then *x* realizes *z*. The same, they note, is true if ‘realization’ is understood in terms of a determination relation (Ibid, n. 16). Fair enough. It is this transitivity, however, that lends support to the idea that what is causally explanatory, ultimately, is what lies at the bottom of the chain of realizations. When there is no unity at the *x* level, it can seem as if the only recourse for a science of conscious experience, the *z* level, is not neuroscience but a science that sticks to the “behavioural or conceptual” levels. Aside from the danger that this outcome forfeits the nociception/pain distinction, it is based on a further assumption that ‘realization’ is the right term to characterise the relationship between different levels of scientific explanation. This is far from obvious. Although Aizawa and Gillett’s paper purports to be about ‘realization’, it is really about ‘compositionality’, which they freely admit is ambiguous. Following Simons (1987), Aizawa and Gillett (2009), p. 182, n.2) use ‘composition’ to encompass a number of relational concepts featured in mechanistic explanations, including ‘realization’ (of properties), ‘constitution’ (of individuals), ‘implementation’ (of processes), and ‘comprising’ relations (e.g., of one power by another). Realization is only one relation among many, not all of which are transitive.

Consider the analysis of rhythmic locomotion offered above. The CPG is *realized* by different sets of neurons in different species, and it *implements* the

¹⁰ Allen (2013) contains, however, a comprehensive overview of the sophisticated learning behaviours of various species of fish and raises the important question which types of learning are consciously mediated.

neurocomputations necessary for rhythmic locomotion, but it's a mistake to say that the set of neurons in each case *implements* the neurocomputations for locomotion—at least they won't in any old configuration! While the CPG is necessary for rhythmic locomotion, it's wrong to say that a determinate set of neurons is necessary, because the sets of neurons constituting the CPG of *Dendronotus* and *Melibe*, while sufficient in each case for the existence of a functioning CPG, are not necessary for rhythmic locomotion. Only the simplified circuit is necessary. Formally, that x (set of neurons) is sufficient for y (circuit) and y is necessary for z (rhythmic locomotion), does not entail that x is necessary for z . In this context, transitivity fails and the explanatory feature is not the one doing the 'realizing' as such, but the one 'implementing' the neurocomputation. By parity of reasoning, techniques enabling us to better understand the neural circuitry performing a psychological function like pain would be more than fully justified.

There is no need, therefore, to resort to "local" or species-specific reductions (cf., Kim 1998, pp. 19–20) either, and as there can be intraspecific variation, this will not help. Nor is it helpful to suppose that because the notion of 'structure' used in neuroscience abstracts from differences that obtain at the neuronal level, it denotes an "ideal type" (Klein 2008). If there is a common neural architecture across distinct species that unifies the type we refer to as 'pain', then while this common architecture is defined in abstraction from what happens at the level of individual neurons, it does not follow that the circuitry so defined is an abstract entity. It would involve actual neural circuitry that does real work controlling the survival-enhancing behaviour of the organism that possesses it. And were we to lack evidence of that circuitry in a given species, we would be able to reasonably infer that the capacity for pain does not there exist.

6 Conclusion

We stress again that we have not tried to establish here that the conclusion of FISH is true, but only to establish the kinds of evidential and epistemic conditions that if met, would allow the argument to be cogent. If the right epistemic conditions *are* met, an absence of evidence argument shifts the burden of proof onto those asserting that a given species is capable of feeling pain to provide evidence that disconfirms the null hypothesis that it is not. FISH* uses modus tollens (FISH**, probabilistic modus tollens) to conclude that it is not likely that fish have the capacity to feel pain.¹¹ Attention to the structure that explains the function of pain in humans is what provides the foundation for the projectability of neurological criteria for pain onto other species. This 'partial reduction' should be enough; no 'global reduction'

¹¹ We have also not attempted here to address the moral question whether it is obligatory to exercise the precautionary principle in regard to non-human species that exhibit nocifensive behaviours (cf: Jones 2016). We recognize the importance of thinking broadly about moral risks when judging whether a species is or is not capable of pain, and acknowledge that the extension of 'harm' exceeds that of 'pain'. But knowing what kind of harm we are dealing with is crucial to knowing how best to respond and exercise our duty of care appropriately. We might, for example, have environmental, ecological or independent moral reasons not to fish certain species, not to fish at certain times or places, or not to fish at all, without needing to equate harms and pains.

specifying necessary *and* sufficient conditions for pain would be necessary or helpful for either explanatory or predictive purposes. Whether FISH or its variants are indeed sound is, however, a matter for empirical interrogation. We have not endeavored, therefore, to show how ‘the explanatory gap’ in the science of pain can be closed, but simply to move the debate about the prospects for a unified science of pain a little closer towards consensus.¹²

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adamo, S. (2018). Sentience, the final frontier. *Animal Sentience*, 21, 3.
- Aizawa, K., & Gillett, C. (2009). The multiple realization of psychological and other properties in the Sciences. *Mind and Language*, 24(2), 181–208.
- Allen, C. (2013). Fish cognition and consciousness. *The Journal of Agricultural and Environmental Ethics*, 26, 25–39.
- Anand, K. J., & Craig, K. D. (1996). New perspectives on the definition of pain. *Pain*, 67, 3–6.
- Armstrong, D. (1968/1993). *A materialist theory of the mind*. London: Routledge.
- Baluška, F., & Reber, A. (2019). Sentience and consciousness in single cells: How the first minds emerged in unicellular species. *BioEssays*, 41(3), 1800229.
- Barberis, S. D. (2017). Multiple realization, levels and mechanisms. *Teorema: Revista Internacional de Filosofía*, 36(2), 53–68.
- Baumgärtner, U., Tiede, W., Treede, R. D., & Craig, A. D. (2006). Laser-evoked potentials are graded and somatotopically organized anteroposteriorly in the opercularinsular cortex of anesthetized monkeys. *Journal of Neurophysiology*, 96, 2802–2808.
- Bayne, T., Howhy, J., & Owen, A. (2016). Are there levels of consciousness? *Trends in Cognitive Sciences*, 20(6), 405–413.
- Bechtel, W., & Mundale, J. (1999). Multiple realizability revisited: Linking cognitive and neural states. *Philosophy of Science*, 66, 175–207.
- Bickle, J. (2003). *Philosophy and neuroscience: A ruthlessly reductive account*. Boston: Kluwer.
- Block, N. (1980). Troubles with Functionalism. In N. Block (Ed.), *Readings in the philosophy of psychology* (Vol. 1, pp. 268–305). Cambridge, MA: Harvard University Press.
- Block, N. (1990). Can the mind change the world? In G. Boolos (Ed.), *Meaning and method*. Cambridge: Cambridge University Press.
- Bock, W. J., & Von Wahlert, G. (1965). Adaptation and the form-function complex. *Evolution*, 19, 269–299.
- Borst, A., & Helmstaedter, M. (2015). Common circuit design in fly and mammalian motion vision. *Nature Neuroscience*, 18, 1067–1076.

¹² We are very grateful for the astute comments we received from two anonymous referees and the editor of *Synthese*. We also want to thank Rod Girle for stimulating discussions about absence of evidence arguments and the audience of the Blackheath Forum for an excellent public discussion in a beautiful setting. Funding for this research was supported by an Australian Research Council (ARC) Discovery Grant (DP200102909; *Toward Closure on the Animal Pain Debate*). We acknowledge the support of the ARC and the University of Queensland in enabling this research.

- Broom, D. M. (2016). Fish brains and behaviour indicate capacity for feeling pain. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 3, 4.
- Brown, C. (2016a). Comparative evolutionary approach to pain perception in fishes. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 3, 5.
- Brown, C. (2016b). Fish pain: An inconvenient truth. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 3, 32.
- Brown, D., & Key, B. (2019). You look but do not find: Why the absence of evidence can be a useful thing. *The Conversation*. <https://theconversation.com/you-look-but-do-not-find-why-the-absence-of-evidence-can-be-a-useful-thing-114988>.
- Copi, I. (1982). *An introduction to logic* (6th ed.). New York: MacMillan.
- Craver, C. (2007). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*. Oxford: Clarendon Press.
- Damasio, A., Damasio, H., & Tranel, D. (2012). Persistence of feelings and sentience after bilateral damage of the insula. *Cerebral Cortex*, 23, 833–846. <https://doi.org/10.1093/cercor/bhs077>.
- De Cornulier, B. (1988). “Knowing whether”, “knowing who”, and epistemic closure. In M. Meyer (Ed.), *Questions and questioning* (pp. 182–192). New York: De Gruyter.
- Demski, L. S. (2013). The pallium and mind/behavior relationships in teleost fishes. *Brain, Behavior and Evolution*, 82, 31–44.
- Denis, D. J., Marouf, R., Rainville, P., Bouthillier, A., & Nguyen, D. K. (2016). Effects of insular stimulation on thermal nociception. *European Journal of Pain*, 20(5), 800–810.
- Derbyshire, S. W. (2016). Fish lack the brains and the psychology for pain. *Animal Sentience*, 3, 18.
- Dimitrijević, M. R., & Nathan, P. W. (1968). Studies of spasticity in man: 3. Analysis of reflex activity evoked by noxious cutaneous stimulation. *Brain*, 91, 349–368.
- Dinet, V. (2016). No cortex no cry. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 3, 7.
- Elwood, R. W. (2017). Assessing negative and positive evidence for animal pain. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 16, 21.
- Favela, L. H. (2017). Mental representations are not necessary for fish consciousness. *Animal Sentience*, 13, 9.
- Gamez, D. (2018). *Fish consciousness*. *Animal Sentience*, 13, 17.
- Garcia-Larrea, L., & Mauguière, F. (2018). Pain syndromes and the parietal lobe. In G. Vallar and H.B.Coslett (eds.), *Handbook of clinical neurology, The Parietal Lobe* (Vol. 151, pp. 207–223). Amsterdam: Elsevier.
- Gerlai, R. (2017). Learning, memory, cognition, and the question of sentience in fish. *Animal Sentience*, 13, 18.
- Geuter, S., Reynolds Losin, E. A., Roy, M., Atlas, L. Y., Schmidt, L., Krishnan, A., et al. (2020). Multiple brain networks mediating stimulus-pain relationships in humans. *Cerebral Cortex*, 30(7), 4204–4219.
- Gillett, C. (2002). The dimensions of realization. *Analysis*, 62, 316–323.
- Gillett, C. (2003). The metaphysics of realization, multiple realizability and the special sciences. *Journal of Philosophy*, 100, 591–603.
- Gillett, C. (2007). Understanding the new reductionism: The metaphysics of Science and compositional reduction. *Journal of Philosophy*, 104(4), 193–216.
- Godfrey-Smith, P. (2016). Pain in parallel. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 3, 21.
- Goldberg, S. (2011). If it were true, I would have heard about it by now. In A. Goldman & D. Whitcomb (Eds.), *The oxford handbook of social epistemology* (pp. 92–108). Oxford: Oxford University Press.
- Gutfreund, Y. (2017). Animal sentience? *Neuroscience has no answers*. *Animal Sentience*, 13, 4.
- Haikonen, P. O. (2016). On the sentience of fish. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 3, 8.
- Harris-Warrick, R., & Ramirez, J.-M. (2017). Neural networks for the generation of rhythmic motor behaviors. In S.L. Hooper & A. Büschges (eds.), *Neurobiology of Motor Control: Fundamental Concepts and New Directions* (pp. 225–262). Wiley.
- Hiddleston, E. (2011). Second-order properties and three varieties of functionalism. *Philosophical Studies*, 153(3), 397–415.
- Huntingford, F. A., Adams, C., Braithwaite, V. A., Kadri, S., Pottinger, T. G., Sandøe, P., et al. (2006). Current issues in fish welfare. *Journal of Fish Biology*, 68, 332–372.
- Jackson, F. (1982). Epiphenomenal qualia. *The Philosophical Quarterly*, 32, 127–136.

- Jones, R. C. (2016). Fish sentience and the precautionary principle. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 3, 10.
- Katz, P.S. (2016) Evolution of central pattern generators and rhythmic behaviours. *Phil. Tran. Roy. Soc. B: Biol. Sci.*, 371:20150057.
- Key, B. (2015). Fish do not feel pain and its implications for understanding phenomenal consciousness. *Biology and Philosophy*, 30, 149–165.
- Key, B. (2016). Why fish do not feel pain. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 3, 1.
- Key, B., & Brown, D. (2018). Designing brains for pain: Human to mollusc. *Frontiers in Physiology*, 9, 1027.
- Kiehn, O. (2016). Decoding the organization of spinal circuits that control locomotion. *Nature Reviews Neuroscience*, 17, 224–238.
- Kim, J. (1992). Multiple realization and the metaphysics of reduction. *Philosophy and Phenomenological Research*, 52, 1–26.
- Kim, J. (1993). Multiple realization and the metaphysics of reduction. In J. Kim (Ed.), *Supervenience and mind* (pp. 309–335). Cambridge: Cambridge University Press.
- Kim, J. (1998). *Mind in a physical world*. Cambridge: MIT Press.
- Kim, J. (2005). *Physicalism, or something near enough*. Princeton, NJ: Princeton University Press.
- Klein, C. (2008). An ideal solution to disputes about multiply realized kinds. *Philosophical Studies*, 140, 161–177.
- Klein, C., & Barron, A. B. (2016). Insects have the capacity for subjective experience. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 9, 1.
- LeDoux, J. E. (2012). Rethinking the emotional brain. *Neuron*, 73(4), 653–676.
- LeDoux, J. E. (2013). The slippery slope of fear. *Trends in Cognitive Sciences*, 17(4), 155–156.
- LeDoux, J. E. (2014a). Coming to terms with fear. *Proceedings of the National Academy of Sciences*, 111(8), 2871–2878.
- LeDoux, J. E. (2014b). Low roads and higher order thoughts in emotion. *Cortex*, 59, 214–215.
- Lenoir, C., Algoet, M., & Mouraux, A. (2018). Deep continuous theta burst stimulation of the operculo-insular cortex selectively affects A δ -fibre heat pain. *The Journal of physiology*, 596(19), 4767–4787.
- Lewis, D. (1970). How to define theoretical terms. *The Journal of Philosophy*, 67(13), 427–446.
- Lewis, D. (1972). Psychophysical and theoretical identifications. *Australasian Journal of Philosophy*, 50(3), 249–258.
- Lewis, D. (1980). Mad pain and Martian pain. In B. Ned (Ed.), *Readings in the philosophy of psychology* (pp. 216–222). Cambridge: Harvard University Press.
- Lindhe Norberg, U. M. (2002). Structure, form, and function of flight in engineering and the living world. *Journal of Morphology*, 252, 52–81.
- Loeser, J. D., & Melzack, R. (1999). Pain: an overview. *The Lancet*, 353, 1607–1609.
- Lovett-Barron, M., Chen, R., Bradbury, S., Andalman, A. S., Wagle, M., Guo, S., & Deisseroth, K. (2019). Multiple overlapping hypothalamus-brainstem circuits drive rapid threat avoidance. *bioRxiv* 745075.
- Maihöfner, C., & Kaltenhäuser, M. (2009). Quality discrimination for noxious stimuli in secondary somatosensory cortex: A MEG-study. *European Journal of Pain*, 13, 1048–e1.
- Manzotti, R. (2016). No evidence that pain is painful neural process. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 3, 11.
- Manzotti, R. (2018). The human nervous system is not the gold standard for pain. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 3, 11.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco, W.H: Freeman.
- Mazzola, L., Mauguière, F., & Isnard, J. (2019). Functional mapping of the human insula: Data from electrical stimulations. *Revue neurologique*, 175(3), 150–156.
- McLaughlin, B. (2006). Is role-functionalism committed to epiphenomenalism? *Consciousness Studies*, 13(1–2), 39–66.
- Merker, B. (2016). Drawing the line on pain. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 3, 23.
- Michel, M. (2019). Fish and microchips: On fish pain and multiple realization. *Philosophical Studies*, 176, 2411–2428.
- Moult, P. R., Cottrell, G. A., & Li, W. C. (2013). Fast silencing reveals a lost role for reciprocal inhibition in locomotion. *Neuron*, 77, 129–140.

- Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, 83, 435–450.
- National Research Council. (2009). *Recognition and alleviation of pain in laboratory animals*. Washington, DC: The National Academies Press.
- Nelson, R. J. (1969). Behaviorism is false. *Journal of Philosophy*, 66(14), 417–452.
- Ng, Y.-K. (2016). Could fish feel pain? A case for a wider, multi-faceted perspective. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 3, 13.
- Parvizi, J., Arciniegas, D. B., Bernardini, G. L., Hoffmann, M. W., Mohr, J. P., Rapoport, M. J., et al. (2006). Diagnosis and management of pathological laughter and crying. *Mayo Clinic Proceedings*, 81, 1482–1486.
- Pedersen, N. J. L. L., & Kallestrup, J. (2013). The epistemology of absence-based inference. *Synthese*, 190, 2573–2593.
- Piccinini, G., & Maley, C. J. (2014). The metaphysics of mind and the multiple sources of multiple realizability. In M. Sprevak & J. Kallestrup (Eds.), *New waves in the philosophy of mind* (pp. 125–152). London: Palgrave Macmillan.
- Polger, T. (2004). *Natural minds*. Cambridge, MA: MIT Press.
- Polger, T. (2009). Evaluating the evidence for multiple realization. *Synthese*, 167, 457–472.
- Polger, T. (2015). Realisation and multiple realisation, chicken and egg. *European Journal of Philosophy*, 23(4), 862–877.
- Polger, T., & Shapiro, L. (2016). *The multiple realization book*. Oxford: Oxford University Press.
- Porro, C. A., Cettolo, V., Francescato, M. P., & Baraldi, P. (1998). Temporal and intensity coding of pain in human cortex. *The Journal of Neurophysiology*, 80, 3312–3320.
- Portavella, M., Torres, B., & Salas, C. (2004). Avoidance response in goldfish: emotional and temporal involvement of medial and lateral telencephalic pallium. *Journal of Neuroscience*, 24, 2335–2342.
- Putnam, H. (1967). Psychological predicates. In W. H. Capitan & D. D. Merrill (Eds.), *Art, mind and religion* (pp. 37–48). Pittsburgh, PA: University of Pittsburgh Press.
- Rainville, P., Duncan, G. H., Price, D. D., Carrier, B., & Bushnell, M. C. (1997). Pain affect encoded in human anterior cingulate but not somatosensory cortex. *Science*, 277, 968–971.
- Rose, J. D., Arlinghaus, R., Cooke, S. J., Diggles, B. K., Sawynok, W., Stevens, E. D., et al. (2014). Can fish really feel pain? *Fish and Fisheries*, 15, 97–133.
- Sakurai, A., & Katz, P. S. (2019). Command or obey? Homologous neurons differ in hierarchical position for the generation of homologous behaviors. *Journal of Neuroscience*, 39, 6460–6471.
- Sakurai, A., Newcomb, J. M., Lillvis, J. L., & Katz, P. S. (2011). Different roles for homologous interneurons in species exhibiting similar rhythmic behaviors. *Current Biology*, 21, 1036–1043.
- Seth, A. K. (2016). Why fish pain cannot and should not be ruled out. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 3, 14.
- Shapiro, L. (2000). Multiple realizations. *Journal of Philosophy*, 97, 635–654.
- Shapiro, L. (2004). *The mind incarnate*. Cambridge, MA: MIT Press.
- Sherrington, C. (1947). *The integrative action of the nervous system*. Cambridge: Cambridge University Press.
- Shoemaker, S. (1984). *Causality and properties identity, cause, and mind* (expanded ed., pp. 206–233). Oxford: Oxford University Press.
- Shoemaker, S. (1998). *Causal and metaphysical necessity. Identity, Cause, And Mind (expanded edition)* (pp. 407–426). Oxford: Oxford University Press.
- Shoemaker, S. (2001). Realization and mental causation. In C. Gillett & B. Loewer (Eds.), *Physicalism and its discontents* (pp. 74–98). Cambridge: Cambridge University Press.
- Simons, P. (1987). *Parts: A study in ontology*. Oxford: Oxford University Press.
- Sneddon, L. U. (2015). Pain in aquatic animals. *Journal of Experimental Biology*, 218, 967–976.
- Sneddon, L. U. (2017). Pain in laboratory animals: A possible confounding factor? *Alternatives to Laboratory Animals*, 45, 161–164.
- Sobel, H. (2009). Modus ponens and modus tollens for conditional probabilities, and updating on uncertain probabilities. *Theory Decis*, 66, 103–148.
- Sober, E. (2009). Absence of evidence and evidence of absence: Evidential transitivity in connection with fossils, fishing, fine-tuning and firing squads. *Philosophical Studies*, 143, 63–90.
- Stauffer, J. R., Jr. (2017). Similar anatomy does not imply comparable function. *Animal Sentience*, 13, 14.
- Strausfeld, N. J. (2005). The evolution of crustacean and insect optic lobes and the origins of chiasmata. *Arthr. Struct. Develop.*, 34, 235–256.
- Strevens, M. (2009). Objective evidence and absence: Comment on Sober. *Philosophical Studies*, 143, 91–100.

- Striedter, G. (2016). Lack of neocortex does not imply fish cannot feel pain. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 3, 15.
- Strother, J. A., Wu, S. T., Wong, A. M., Nern, A., Rogers, E. M., Le, J. Q., et al. (2017). The emergence of directional selectivity in the visual motion pathway of *Drosophila*. *Neuron*, 94, 68–182.
- Tracey, I. (2005). Nociceptive processing in the human brain. *Current Opinion in Neurobiology*, 15(4), 478–487.
- Wager, T.D., & Barrett, L.F. (2017). From affect to control: Functional specialization of the insula in motivation and regulation. *BioRxiv* 102368.
- Wagner, C. G. (2004). Modus tollens probabilized. *The British Journal for the Philosophy of Science*, 55, 747–753.
- Wake, M. H. (1992). Morphology, the study of form and function, in modern evolutionary biology. *Oxford Surveys in Evolutionary Biology*, 8, 289.
- Walton, D. (1996). *Arguments from ignorance*. University Park, PA: Pennsylvania State University Press.
- Wee, C. L., Nikitchenko, M., Wang, W. C., Luks-Morgan, S. J., Song, E., Gagnon, J. A., et al. (2019). Zebrafish oxytocin neurons drive nocifensive behavior via brainstem premotor targets. *Nature Neuroscience*, 22, 1477–1492.
- Williams, A. C. D. C., & Craig, K. D. (2016). Updating the definition of pain. *Pain*, 157, 2420–2423.
- Wimsatt, W. (1994). The ontology of complex systems: Levels, perspectives and causal thickets. *Canadian Journal of Philosophy*, 20, 207–274.
- Wimsatt, W. (2007). *Re-engineering philosophy for limited beings: Piecewise approximations to reality*. Cambridge, MA: Harvard University Press.
- Woodruff, M. L. (2017). Consciousness in teleosts: There is something it feels like to be a fish. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 13, 1.
- Woodruff, M. L. (2018). Pain in fish: Evidence from peripheral nociceptors to pallial processing. *Animal Sentience: An Interdisciplinary Journal on Animal Feeling*, 21, 2.
- Woodworth, R. S., & Sherrington, C. S. (1904). A pseudoaffective reflex and its spinal path. *Journal of Physiology*, 31, 234–243.
- Zeki, S., & Bartels, A. (1998). The autonomy of the visual systems and the modularity of conscious vision. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 353, 1911–1914.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.