

What is a pain in a body part?

[Murat Aydede](#)

Department of Philosophy
University of British Columbia
maydede@mail.ubc.ca

Abstract. The IASP definition of 'pain' defines pain as a subjective experience. The *Note* accompanying the definition emphasizes that as such pains are not to be identified with objective conditions of body parts (such as actual or potential tissue damage). Nevertheless, it goes on to state that a pain "is unquestionably a sensation in a part or parts of the body, but it is also always unpleasant and therefore also an emotional experience." This generates a puzzle that philosophers have been well familiar with: how to understand our utterances and judgments attributing pain to body parts. (The puzzle is of course general extending to all sensations routinely located in body parts.) This work tackles this puzzle. I go over various options specifying the truth-conditions for pain attributing judgments, and at the end make my own recommendation which is an adverbialist, qualia-friendly proposal with completely naturalistic credentials that is also compatible with forms of weak intentionalism. The results are generalizable to other bodily sensations and can be used to argue, quite generally, for a qualia-friendly adverbialist (but naturalist) account of perception.

1 Preamble

Pain scientists and clinicians often say that pain is a subjective experience. This is epitomized in the widespread acceptance and use of the IASP definition of pain:¹

Pain is “an unpleasant sensory and emotional experience associated with actual or potential tissue damage, or described in terms of such damage.” (IASP Taxonomy 1979/2011)

The *Note* to the definition adds that “Pain is always subjective,” and goes on to state that:

“... This definition avoids tying pain to the stimulus. Activity induced in the nociceptor and nociceptive pathways by a noxious stimulus is not pain, which is

¹ The International Association for the Study of Pain (IASP) now keeps its official list of terminology and definitions online: <<http://www.iasp-pain.org/Education/Content.aspx?ItemNumber=1698>>

always a psychological state, even though we may well appreciate that pain most often has a proximate physical cause.” (*ibid.*)

Even those who take issue with the IASP definition typically do not question the claim that pains are subjective experiences.

Given that toes, buttocks, guts and other bodily parts and organs aren't regions of the body that can have experiences in them, the subjective experience conception of pain creates a *tension* with people's routine and otherwise unremarkable practice of attributing pains to bodily parts like “I have a pain in my toe,” “my back hurts,” “I have a toothache,” etc. One might wonder whether the scientists and clinicians may be advising against the practice of locating pains in body parts. But of course, they themselves routinely engage in this practice. Are they aware of the tension? They seem not. The very same *Note* contains the following sentence: “[pain] is unquestionably a sensation in a part or parts of the body, but it is also always unpleasant and therefore also an emotional experience” (*ibid.*)

The IASP definition and the remark in the accompanying *Note* seem to rule out identifying pain with any objective condition of the body part to which a pain is attributed, such as any actual or potential tissue damage, injury, or any kind of physical disturbance. The implication seems to be that although a pain is typically associated with such a condition, it's not to be identified with such. Can the pain scientists or clinicians be serious in affirming the implication that when one sincerely claims to have a pain, say, in one's toe, one is attributing a subjective experience (something like an emotion) to one's toe? Can they be serious in thinking that people can have subjective experiences in body parts like toes, guts, buttocks, etc.? This seems very unlikely.² Subjective experiences like pains are said to be realized in the brain (or the central nervous system). Indeed, many pain scientists think of pain experiences as being in the brain. Clearly, toes, guts, buttocks and the like don't have the complex functional organization sufficiently similar to the brain (any brain) to support subjective experiences.

So, we seem to have a puzzle. This puzzle isn't confined only to the conception of pain that scientists and clinicians have. It seems widespread in the larger population reflecting the puzzle embedded in the folk conception of pain.³ To explicitly state: the puzzle stems from the fact that we seem to conceive of pains as both subjective experiences and as locatable conditions of the very objective bodily parts and organs. The folk may be excused for being sloppy, inattentive, and sometimes even confused. This wouldn't be very surprising. The folk aren't doing any kind of *research* relying on this conception after all — they are most often just trying to get by in their hectic daily lives.

² Indeed, the preliminary results of an empirical survey among scientists and clinicians strongly suggest that they don't think the relevant body parts can have subjective experiences in them. See below.

³ In the philosophical literature, this puzzle has come to be called the “paradox of pain” (Hill 2006, 2009; Borg et al. forthcoming; Bradley, 2018). I prefer ‘puzzle’ over ‘paradox’ since the latter may imply that the conception is internally incoherent. I don't want to make this stronger claim (and I have never made such a strong claim in print or otherwise previously). In my view, there is only an internal conceptual tension, a puzzle, not a formal incoherence — see my (2005/2013) and below.

But the puzzle becomes more puzzling when it's realized that pain scientists and clinicians worked long and hard to come up with the IASP definition and the accompanying *Note*.⁴ Pain science is the science of pain so defined. Pain clinicians treat patients with pain so defined. How can the pain scientists and clinicians be confused? If there is a confusion, it seems to be of such an elementary and fundamental kind that makes one wonder how the scientists and clinicians can seem to be unaware of it. Indeed, if there is a confusion, it doesn't seem to be even below the surface (to wit: "[pain] is unquestionably a sensation in a part or parts of the body, but it is also always unpleasant and therefore also an emotional experience" *ibid*).⁵ Also, if there is a confusion, wouldn't one expect this to somehow show up somewhere — perhaps having a somewhat adverse impact on scientific research or clinical practice? But as far as anybody knows, there seems to be no trouble or difficulty in any aspects of scientific pain research or clinical practice that is traceable to this "confusion."

So we seem to have a puzzle of a very peculiar sort. Scientists, clinicians, and the folk alike seem to be utterly competent in attributing pains to body parts while talking of pains as subjective experiences without much hesitation in their discourse or in their practical dealings with each other as well as in their clinical or research settings. Everybody seems to implicitly know what the truth-conditions of pain attributing utterances are so they routinely speak and evaluate the relevant language without any signs of confusion — even though most seem to draw a blank when asked to articulate what these truth-conditions are.

This puzzling situation calls for an explanation. In particular, we need to understand how scientists and clinicians can be operating with two *apparently* conflicting conceptions of pain without any discomfort or sign of confusion, especially when the apparent conflict seems so manifest. My own view is that despite appearances, these two conceptions of pain, call them *bodily* versus *experiential* conceptions, are not actually in conflict, and consequently, there is no serious or deep confusion in the conception of pain had by scientists, clinicians, and the folk alike. In what follows, I will try to bring this out by going over the options about the truth-conditions of utterances of pain to body parts, and at the end (§3), offer my own proposal, which will have wider implications in the philosophy of perception.

2 Options

Consider John who has started to feel a pain in his toe after a long weekend run. John is a mature person with an intact healthy body. He sincerely makes the following utterance:

⁴ See Merskey 2008 and Aydede 2017c for critical discussions of the intellectual history and the assumptions behind the IASP definition.

⁵ The problem isn't of course confined to pains but it's about all similar bodily sensations, such as itches, tickles, pins-and-needles sensations, etc. We routinely locate sensations in parts of our bodies and mean these sensations to have an awareness-dependent existence. Note how natural the quoted sentence from the *Note* sounds — but a moment's reflection should make it remarkably puzzling ("an emotional experience in one's toe!").

“I have a pain in my toe.”

What makes this utterance true? We would like to know the truth-makers of such utterances locating pains in body parts (call these, *L-utterances*). More specifically, we want to know both the logically necessary *and* sufficient conditions for L-utterances to be true. What are the truth-conditions of L-utterances?

We already ruled out this option above:

(1) John’s toe has a “painful” subjective experience in it.

Given the remarks in the *Note* added to the IASP definition, quoted above, we can rule out the following option too:

(2) John’s toe has some kind of physical injury (disorder, tissue damage) in it.

The appended *Note* warns us against tying the pain to physical noxious stimuli and explicitly disapproves of identifying pain with the activity of nociceptors or nociceptive activity in the afferent pathways. It would be particularly weird if it left open the option of identifying pain with the actual or impending tissue damage that is typically in between the noxious stimuli causing it and the nociceptor activity caused by it.

But perhaps what makes John’s utterance true is the fact that:⁶

(3) John’s toe has some kind of tissue damage in it, *and* John feels it in a “painful” way.

This option does better: (3) seems sufficient for John’s utterance to be true. But (3) is clearly not necessary. Suppose that John’s pain in his toe is due to a slightly compressed and inflamed nerve in his lower back that started to act more aggressively after his long run. There is nothing physically wrong with his toe — no actual or potential tissue damage of any kind detectable by nociceptors. The pain John feels in his toe is now what pain scientists call a “referred pain.” Under such a scenario, John’s utterance remains true but (3) becomes false. Therefore, (3) cannot capture what makes John’s utterance true.

It is important to elaborate this point. Referred pains are no less genuine pains. And they are not mislocated pains either. John doesn’t feel the pain in his lower back, he feels the pain in his toe. It is his toe that has pain in it. When asked whether he has any pain in his lower back, John sincerely reports that he has none. John is not confused. Nor is any physician who agrees with John confused. The physical disorder that is the cause of

⁶ Here and throughout this essay, I’ll be using ‘painful’ to mean ‘of the phenomenal character of sensory pain.’ Usually ‘painful’ is used to pick out the affective hurting aspect of pain in the literature. Although my use doesn’t necessarily exclude this aspect in this paper, I won’t have this meaning in mind here. “Painy” would have been a better option, but it’s not an established English word.

John's pain is in his lower back. So to this extent, it may be plausible to say that there is something misleading about John's *pain in the toe* if we expect pains in body parts to *signal* physical disorders in those parts. But when we don't speak loosely and are careful, we don't stop believing that John has a pain in his toe upon learning where the cause of the trouble is if John continues to feel the pain in his toe. So (3) won't do.

To say that John is having a referred pain in his toe is to say that John has a pain in toe and the cause of his pain is somewhere other than in his toe. So, John still has a pain in his toe and we are trying to understand what makes this true. At this point it seems clear that John's pain in his toe has nothing to do with any physical or objective condition of his toe. Perhaps we can capture John's utterance with the following:

(4) Whether or not John's toe has any physical disturbance in it, John's toe has a subjective "painful" quality in it, and John now senses or is aware of this quality.

With this option we seem to be in the realm of fantastical metaphysics. It is not quite clear what it means to say that John's toe has a subjective quality in it. Nevertheless, from a phenomenological point of view, it's also clear that it seems to John *that there is* a "painful" hurting quality instantiated in his toe, and *that* he is aware of it. What makes this quality *subjective* is the fact that it also seems to John *that* this quality is awareness-dependent. It would seem to disappear the moment he ceases to be aware of this quality (let's say, *despite his full attention*). John might worry about the *physical* condition of his toe or back that has caused the pain in his toe in the absence of his feeling pain (let's say, after taking strong analgesics), but it seems very unlikely that John would worry about the "current" pain in his toe if he is not feeling it *at all*. (4), then, is an attempt to respect the first-person phenomenology (how things seem to a person from the person's experiential perspective) by turning it into a metaphysical account. On this view, there really *are* subjective qualities literally instantiated in body parts despite their having an awareness-dependent status. An analogy might help here. Consider using your flashlight in the dark to illuminate your environment. The circular bright spot falling on objects around you is dependent on whether you shine the light on them with your flashlight. The light spot really exists on objects when you illuminate them with your flashlight. The bright spot on objects is one thing, and your act of shining light on them with your flashlight is another. Similarly, John's act of feeling the painful quality is one thing, and the quality instantiated in John's toe is another — even though the quality's instantiation in John's toe metaphysically depends on his feeling it.⁷

Unlike the spotlight, however, the subjective quality in Joe's toe has further peculiarities. Not only is it subjective, but it also seems epistemically private in the following sense: there seems to be a radical asymmetry between the way John has epistemic access to it and the way others may have access to John's pain in his toe. This

⁷ This analogy shouldn't be taken *very* seriously. It will break down quickly once we realize that the causal processes in the two cases are very different. Nevertheless, it helps with the general intuitive picture about how (4) is meant to be taken.

radical difference in epistemic access gives rise to another radical difference in epistemic and practical/prudential authority John has over the pain in his toe over the kind of authority others might have about John's pain.

Despite its near-perfect match with first-person phenomenology, what makes (4) difficult to accept is its apparently fantastical nature as a piece of metaphysics.⁸ Subjectivity, privacy, being a source of privileged and perhaps incorrigible knowledge, all seem to be mysterious features associated with either sense-data or sensory experiences with qualia — not associated with body parts. Sense-data are thought to be metaphysically dubious *entities*, whereas qualia are thought to be metaphysically dubious *qualities* of sensory experiences for similar reasons. If there is a way of making (4) metaphysically acceptable for both naturalistically oriented philosophers and pain scientists, then (4) would seem to fit the bill well for what it is that makes John's utterance true. Later on, I'll have a proposal that might accommodate most of its attractiveness. But for the moment we may want to look into another option that may be suggested by the fact that subjectivity, privacy, etc., all seem to be features associated with sensory and emotional *experiences* and *their* qualities.

Perhaps we can turn the experiential conception of pain into an account that will answer the bodily conception as well.

(5) Whether or not John's toe has any physical disturbance in it, John has a subjective experience that attributes to his toe some sort of "painful" quality.

There are different ways of understanding this that we need to carefully sort out. The most natural interpretation is to take the experience to *represent* a "painful" quality as being in John's toe. The attribution involved isn't as if the experience is throwing a quality at the toe so that the toe can catch it — it is an intentional attribution. If so, the experience has *accuracy conditions*: it "says" of the toe that it has a certain quality in it. What it says is correct if and only if the toe has the quality attributed — false or incorrect otherwise. We now have a similar worry about what this represented/attributed quality is.

Perceptualists and representationalists (PR theorists)⁹ claim this quality to be an objective/physical condition of the toe, such as actual or potential tissue damage or injury (disorder, for short). On this view, what makes John's utterance true is the fact that John undergoes a subjective experience that (non-conceptually or in an analog format) represents his toe as having a physical disorder of the appropriate kind in it. This experience is veridical if John's toe has indeed such a disorder in it (and is relevantly caused by the disorder) — non-veridical otherwise. So if John's pain in his toe is caused by a nerve compression in his lower back, then John's subjective experience is inaccurate, non-veridical. But, according to PR theorists, this subjective experience is, because of its

⁸ Is there anybody who actually defends such a view? It's not clear. John Hyman (2003) can probably be interpreted as sympathetic. Here one is tempted to describe the view as primitivism about bodily pains. But awareness-dependence doesn't seem to fit with the general primitivist idea. Adam Bradley (2018) has developed a view that probably comes closest.

⁹ See Bain (2003, 2007), Tye (2006), Hill (2006, 2009).

intentional content, still a genuine pain experience despite its inaccuracy. So, according to PR theorists, pains are primarily subjective experiences, just as the IASP definition says. But they claim that they can make sense of pain locating judgments (*L-judgments* expressed by L-utterances) if the representational content of such experiences is taken into account. What makes John's utterance true is the fact that:

(5i) Whether or not John's toe has any physical disturbance in it, John has a subjective experience that represents his toe as having a physical disturbance in it.

Whether or not John knows it or can articulate it, when John sincerely reports a pain in his toe by uttering "I have a pain in my toe," he is in fact expressing this fact. (5i) articulates the truth-conditions of John's utterance as well as his relevant belief thereby expressed. Another way to understand (5i) is to make a relational property out of the intentional content of the experience: John's toe has a pain in it if and only if John's toe is the intentional target of John's experience attributing a disorder to it. Such properties are sometimes known in the philosophical community as the inverse intentional properties and are subject to debate as to whether they are genuine properties because they are said to flout the Leibniz Law.¹⁰ But we need not worry about this debate here. The point is that attributing an inverse intentional property to the toe seems to make sense of what it is that we are attributing to the toe while at the same time we can trace the subjectivity and privacy involved in this attribution to the subjectivity and privacy of the experience itself, thus apparently removing our initial puzzle.¹¹

This seems promising, but the view has quite unattractive consequences. One is already mentioned and may be considered minor: it makes a lot of genuine pain experiences literally mistaken — a lot of them turn out to be incorrect or seriously in error. We don't normally talk about pain experiences as literally mistaken in the sense that they "say" something false. To be sure, we take them to be sometimes *misleading* given our natural expectation of finding something physically wrong in those body parts where we feel pain. But this is quite different than taking them to be literally mistaken. A PR theorist may find this to be a small price to pay compared to what the view seems to accomplish in bringing a unifying solution to the initial puzzle (the tension between the bodily and experiential conceptions of pain).

But (5i) in the hands of PR theorists generates another puzzle that cries out for explanation. Following the IASP definition, PR theorists consider pain to be primarily a

¹⁰ Consider the following fallacious argument: Mark Twain is known to be the writer of *Huckleberry Finn*; Samuel Clemens is *not* known to be the writer of *Huckleberry Finn*; therefore, Mark Twain isn't one and the same person as Samuel Clemens.

¹¹ (5i) is a version of intentionalism that's come to be known as *evaluativism* (Bain 2013) when supplemented with the claim that the experience further represents the disturbance as bad. This further claim is added to capture the affective-motivational dimension of pain experiences. Klein (2015) developed a purely imperativist version of intentionalism, according to which to have a pain sensation in one's body part is just for one to have been issued (by one's body) a command to protect that part in a certain way. Klein denies that pain experiences have any descriptive content supported by nociception. I have criticized Klein's account elsewhere (2017d).

subjective experience (realized in the central nervous system). Their proposal about how to make sense of pain attributions to body parts is to postulate that pain experiences are representational — they have an intentional content. This content is about *what* happens *where* in the body. Pain experiences are like other sensory experiences in telling us what is happening where so that we can form direct *perceptual beliefs* on their bases. The point of having sensory experiences is that we know how to epistemically and practically deal with our extra-mental (worldly, including bodily) environment. Our ordinary epistemic and practical focus in having sensory experiences as representational media is on *what* they represent to us accurately, not *that* they represent to us *whatever the accuracy*. So if pains are primarily experiences that represent physical disorders in body parts, then one would expect the L-utterances and the beliefs expressed thereby to reflect this fact and therefore come to be true or false depending on whether or not the represented bodily parts have the represented disorder. But what we find instead is quite different, namely, (5i). (5i) is precisely what a PR theorist ought *not* to say as a response to our question about what makes John’s utterance (“I have a pain in my toe”) true.

A PR theorist might reply that L-utterances are not perceptual — they are introspective. Their form is misleading. They are really different ways of making an introspective report to the effect that one is simply having a subjective pain experience with a certain intentional content about body parts and their disorders. They are misleading because pains are really in the brain, they are never in one’s body parts — they can’t be. But they nevertheless represent disorders in body parts. But we are rarely interested in making perceptual judgments expressed by these L-utterances to the effect that we have a *disorder* in a body part. Appearances and expectations to the contrary, our L-utterances don’t report disorders. Rather they report our experiences whose typical *causes* are the bodily disorders that we sometimes report as such. Thus, PR theorists are bound to concede that if we sometimes make “perceptual” judgments about disorders *as such* based on our pain experiences, we do that by *inference* — given our natural expectation about the regularity of such causal connections between (actual or impending) physical injuries and pain experiences (at least under certain circumstances). But this is different from genuine or *direct* perceptual judgments that are not thus inferentially mediated. Thus the admission that L-utterances are not perceptual doesn’t help: *we are* now missing a proper (direct) perceptual judgment on this view. PR theorists make a mystery out of their own claim that pain experiences are perceptual/representational — by conceding that we don’t seem to have (ever) genuine direct perceptual judgments properly based on them. So I will leave (5i) aside.¹²

Our troubles seem to stem from the fact that we wanted to interpret (5) as saying that the quality attributed to body parts is some sort of objective condition (disorder) of those parts. What if we drop this interpretation and opt for something else as the quality attributed to body parts? What would this be?

¹² I’ve criticized perceptual/representational views of pain extensively in the literature. I refer the reader to Aydede (2006, 2009, 2017a,b) and Aydede & Fulkerson (2014).

Given that we are taking the language of ‘attribution’ as intentional in (5), whatever candidate we come up with as the quality attributed, it will generate accuracy conditions for pain experiences. This is in fact the source of our trouble with (5). Our options here are quite limited. This quality is either a naturalistic quality or it is not. If it's a naturalistic quality, I don't see any plausible candidate other than some kind of physical disorder (say, actual or potential tissue damage or injury).¹³ But given the troubles with (5i), we can rule this possibility out. But then what could this quality be such that it can generate accuracy conditions when attributed to body parts? Consider (5ii) where the quality is taken to be a *sui generis* primitive quality not reducible to anything else:

(5ii) Whether or not John's toe has any physical disturbance in it, John has a subjective experience that attributes to his toe some sort of *primitive* “painful” quality.

With this option, John has a pain in his toe if and only if he undergoes a subjective experience that represents his toe as having a primitive "painful" quality. Let's call this *sui generis* quality *P*. Is this a veridical experience? Or in the material mode: does John's toe have this primitive quality in it? Hard to tell. *P* is not a physical quality, not identifiable or even (strongly) correlated with any physical condition of the body (beyond the CNS). So how do we make a judgment of veridicality?

The most plausible way to proceed here is to say that *P* is only a represented quality: the experience presents his toe to John as if it instantiated *P* but in fact *P* may not be instantiated in the actual world — not in the toe anyway. So the experience is likely not veridical. In fact, all pain experiences may be non-veridical according to this interpretation. *P* may live only in the intentional *content* of the pain experience projected on to John's toe as if instantiated there. But *P* may in fact not be instantiated anywhere — including the experience. *P* is *represented* as instantiated, for sure, but in fact is not instantiated anywhere. Such a view is sometimes called figurative projectivism.¹⁴

The important point, however, is that for John's utterance to be true, it doesn't matter whether John's toe has *P* or not. According to option (5ii), the truth-conditions of L-utterances will be fixed independently of whether these experiences are accurate or veridical. In so far as John is undergoing a subjective experience that is presenting his toe as having this "painful" quality, *P*, John is having a pain in his toe. *The primitive quality P and the property of being represented as having P are distinct properties.* It's the former property that is represented by the experience, not the latter — but see below. Bodily pain, John's pain in his toe, is therefore not to be identified with *P*, but rather with the inverse intentional property of being represented as having *P*. We want John's utterance to come out true. So the concept John uses when he makes his report cannot be the mere concept expressing the property of *P* — as expected under this proposal. For his labeling use of

¹³ This will be qualified shortly — see below, fn. 15.

¹⁴ See Shoemaker (1994). A similar view is defended by Adam Pautz (2010). David Chalmers' (2006) views on Edenic content are also similar.

this concept to locate an instance of *P* in his toe would make his utterance false — as *P* is likely nowhere instantiated. With this option, the concept John (along with the rest of us) is utilizing when he makes his true utterance must therefore be a concept expressing something like the inverse intentional property of being represented as having *P*. But this concept is *not* the concept of *P*.

At this point, however, the claim that John's experience *represents* his toe as having this primitive quality becomes unmotivated — for the purposes of understanding L-utterances. The notion of representation not only does no work here, it puts extra burden on the effort to remove the tension between bodily and experiential conceptions of pain by generating accuracy conditions for the subjective experience that the IASP definition identifies pain with. It makes the pain experiences always mistaken (in error), even though, surprisingly, the L-judgments based on these experiences reporting pains in body parts come out true — as specified by (5ii). Note that this is a consequence of the fact that the concept utilized in the L-judgment attributes a different property than represented by the experience itself — hence the costly discrepancy between the contents of the experience and of the L-judgment directly based on it.

There is a more straightforward form of projectivism whose discussion will make the frivolousness of insisting on representationalism even clearer:

(5iii) Whether or not John's toe has any physical disturbance in it, John has a subjective experience with a “painful” quality *P* that attributes itself to John's toe — equivalently: that represents itself as being instantiated in John's toe.

This view is sometimes called literal projectivism (see Shoemaker 1994). The primitive quality *P* is not merely a represented property but it's also literally instantiated by John's pain experience. However, *P* is a self-representational property too: it represents itself as being instantiated somewhere else, namely in John's toe. So the pain experiences are always in error as in figurative projectivism: *P* is not instantiated in John's toe, or in any bodily part (except in the CNS). Nevertheless, on this view, John's utterance still comes out true despite the massive error the experience commits that grounds the L-judgment.¹⁵

Sydney Shoemaker (2000) has advanced an interesting variant of (5iii) for bodily pains. His view about what he calls occurrent appearance properties can be taken to imply that John's utterance is true just in case: John has an experience with a subjective “painful” quality (quale) *P* that represents itself as being occurrently caused by his toe. This view fares better in normal acute pains in delivering veridical experiential contents, but makes most referred pain cases involve misrepresentation. Shoemaker seems to think that “a 'referred pain in the arm' is not really in the arm” (2000: 262) — since the occurrent appearance property represented in the experience is not actually instantiated in the arm: it is in fact the heart that is occurrently causing *P*. To Shoemaker's credit, this is an attempt

¹⁵ Although I am interpreting the quality *P* to be primitive in (5iii), it need not be a metaphysically primitive quality as in (5ii). If one is a naturalist or physicalist like Shoemaker and myself, the quality *P* may be a quale that is realized by physical/functional properties of the neural activity in the brain. Indeed, this is what I will be assuming in my own proposal below when I talk about qualia as predicates.

to make the contents of the experience and of the L-judgment directly based on it semantically aligned. If we take a pain in a body part *b* to be the relation that *b* instantiates by occurrently causing an experience with quality *P* in subject *s*, then John doesn't have a pain in his toe, although he falsely thinks he does. So on Shoemaker's view, in a referred pain case, both the pain experience and the L-judgment based on it are mistaken — have false intentional content. The implication of the view that John doesn't have a pain in his toe is unacceptable.¹⁶ The toe is where John feels the pain, so his L-utterance is true, even when there is nothing physically wrong with the toe and the toe is not causally responsible for John's feeling pain in his toe.

A way of saving Shoemaker's proposal is to make the same move as in the case of other representationalist options by saying that this experiential error involved, by fiat, should make no difference whatsoever in assigning truth-conditions to John's utterance. In other words, it's both necessary and sufficient for John's utterance to be true that he has a subjective experience with a quality *P* that (mis)represents his toe as (occurrently) causing itself (*P*). So the veridicality of pain experiences is irrelevant whatsoever to the truth of those L-judgments directly based on them.

To better see the frivolousness of representationalism in the latter two formulations of (5), we can further note that another way to describe these options is to say that his toe somatosensorially appears to John to have a *subjective* primitive quality: this quality is appearance-dependent even to John. It *is* in fact an appearance property: it exists only in so far as it appears. John won't claim his toe to *have* the pain if there is no appearance of the relevant kind (say, despite John's full attention). It doesn't make much sense to ask whether John's toe might in fact lack the quality despite the appearances. But if so, this makes sense only if we don't take the appearances to be making a semantic claim on reality — on whether or not John's toe has in fact the "represented" quality. In other words, we have the option of not taking the appearances in pain experiences intentionally.

The situation is quite different here than the case where we report how a ripe tomato appears to one by saying "it appears red to me." Here the redness of the tomato is not (merely) an appearance property so we can meaningfully ask whether the tomato is in fact red despite how it appears to me. The appearance here makes a claim on the reality¹⁷ — inviting us to take the appearance intentionally — so that if the tomato turns out not to be red, one can blame my visual experience for providing me with "false appearances."

¹⁶ In pilot experimental studies I have been conducting with Adam Shriver surveying the linguistic intuitions of pain scientists and clinicians, *all* the pain scientists and clinicians (more than 60, except one) have agreed that John *does* have a pain in his toe, even when it is a referred pain completely caused by his lower back. Shoemaker has told me (in correspondence) that he would now agree that John's utterance is true and that he has a pain in his toe even when it's completely a referred pain. I haven't touched on phantom limb pains in this essay as they complicate things (in ways irrelevant to our current discussion) due to reference failure because of missing body parts — see below.

¹⁷ The comparison with colour perception here doesn't depend on any particular views about the ontology of colours — although I am excluding those views that are *entirely* subjective so that our colour attributing judgments are not at all constrained by *any* extra-mental condition of the world. I consider such views to be implausible.

Option (5), then, is not satisfactory no matter how we take the attributed quality to be. The moment we insist on John's experience intentionally *attributing* a quality to his toe, we generate accuracy conditions for the experience that turn out to be irrelevant for the truth of the L-utterance or L-judgment based on it. John's subjective experience doesn't seem to *represent* his toe as having a certain quality in it or as being in a certain appearance-*independent* condition. This is not to say that John's experience is completely representation-free. There is clearly a representational element in so far as John's experience is *about his toe*: it makes *reference* to his toe as part of his body. Rather, it is to say that it doesn't do any *predication* by attributing any property or quality to it. In having the experience John has, his toe appears to John in a certain way, it appears to him "painfully." This *way of appearing* is not intentional in the case of "pain in the toe": it doesn't represent any quality or property as instantiated in the toe. How can we use this insight in our investigation about what makes John's utterance true?

Here is my own proposal that I claim does the best job in relieving the tension between the bodily and the experiential conceptions of pain without generating any error on the part of the experience.

(6i) Whether or not John's toe has any physical disturbance in it, John has a subjective experience that presents his toe in a certain way — "painfully."

Or, equivalently:¹⁸

(6ii) Whether or not John's toe has any physical disturbance in it, John has a subjective experience that makes his toe feel or appear to him in a certain way — "painfully."

3 The coherence and unity of the bodily and experiential conceptions

What follows is a rough outline of how to implement (6) in a naturalist framework that comports well with the core traditional understanding of pains as both locatable items and subjective experiences. The proposal is weakly intentionalist, qualia-friendly, but fully naturalist, which has close affinities with adverbialism in philosophy of perception as is clear from the formulation of (6).¹⁹

The fundamental insight behind (6) is that there is a quasi-intentional structure to the subjective experience, which the IASP definition identifies with pain. Like most (perhaps all) sensory experiences, pain experiences have a syntactic structure with a reference/predicate divide. John's pain in his toe is an experiential presentation of his toe.

¹⁸ Note that these formulations avoid the construction "... presents his toe *to be* a certain way" or "... appear to him *to be* a certain way," which suggests the representationalist interpretation.

¹⁹ For more details of the view, especially about its naturalistic credentials, see my (2017b: §6). The general information-theoretic treatment of sensory experiences feeding into conceptual structures is developed in Aydede & Güzeldere (2005) with an eye to solve the puzzle of qualitative consciousness.

This is the *referential* aspect of his experience. John's pain experience presents his toe to him in a certain way. What way this is is determined by the *predicative* structure of his experience. In genuinely perceptual/sensory experiences as arising for instance from visual, auditory and tactile modalities, the predicative structures function as property-attributive devices and the properties they attribute are the sensible objective features of the (usually) distal stimuli. But the claim here is that the predicative structure of pain experiences does not attribute any properties or qualities to the distal stimuli (bodily parts) they refer to. Unlike the sensory experiences arising in most other modalities, pain experiences are not perceptual, nor are they fully representational.²⁰

The talk of a reference/predicate divide in the structure of sensory experience might be disorienting as the distinction is usually made regarding linguistic or sentential syntactic structures. It need not be: we do have applications of this sort of syntactic structure in non-sentential or non-discursive representational media. As an example, think of a paper marine chart of a certain marine region (say, the Burrard Inlet). The points or regions on the paper chart as fixed by the horizontal/vertical side coordinates of the chart (as labeled in terms of the meridians of latitude and longitude) will pick out or refer to actual points or regions in the inlet. This is the *referential* aspect of the representational medium (paper chart). What colors, lines, or marks there are on those points or regions of the paper chart will then tell us what oceanographic or geographic *properties* (e.g., depth, currents rates or directions, submerged rocks, etc.) are true of the corresponding places in the inlet. This is the predicative aspect. Here we have a representational vehicle (the paper chart) with a syntactic structure roughly corresponding to referential and predicative functions. The claim is that sensory experiences as intentional structures have a similar referential/predicative divide. Despite the presence of a predicative structure, however, the representational format of sensory experiences is not conceptual, similar to the way the marine chart is *not* a discursive or sentential representation.

The idea in (6) is that pain experiences present body parts in a certain way — painfully, in all the phenomenally determinate ways our body parts are presented to us in pain. The qualitative or phenomenal character of pain experiences is determined by their predicative structure.²¹ This view is not only compatible with but naturally invites a qualia-friendly treatment of the subjective qualities of pain experiences. Pain qualia are implemented by the proprietary predicates of the nociceptive system but are not representational in that they don't function to attribute any quality to the body parts presented in the same experience. Pain qualia are the phenomenally determinate *ways* in which our body parts are presented to us or appear to us.

It is instructive to compare (6) to different versions of (5) above. The options in (5) are all representationalist: they claim that John's subjective experience represents his toe as

²⁰ Thus I am not arguing against any kind of representationalism about perceptual experiences in this paper. I am a weak representationalist about standard perceptual modalities — see Aydede & Güzeldere (2005). For a more rigorous and detailed examination of the reference/predicate divide in sensory experiences, see Clark 2000.

²¹ Plus, affect (hedonic value) — to be handled as further (second-order) adverbial modification of the instantiations of sensory predicates. See Aydede (2014, 2018) for details.

instantiating a certain quality *P*. According to (5i), this quality is some sort of physical disorder in the toe. According to (5ii) and (5iii), *P* is a *sui generis* primitive quality represented to be instantiated in John's toe. (5ii) takes *P* to be instantiated nowhere — *P* lives only in the representational content of pain experiences. (5iii) takes *P* to be instantiated in John's pain experience: *P*, we may say, is the pain quale instantiated in the pain experience and determines John's pain phenomenology. This part of (5iii) is compatible with (6).²² But (5iii) makes the further claim that *P* is a self-representational quality: it represents itself as instantiated in John's toe when in fact nothing of this sort can be instantiated in toes or other bodily parts. Thus (5ii) and (5iii) make pain experiences to commit massive error at all times — even though at the end this error doesn't affect the assignment of correct truth-conditions to L-utterances (thus, showing the frivolousness of representationalism). (6) is like (5iii), but is not (fully) representationalist. There is a pain quale instantiated in the experience which presents John's toe to John in a certain way: this is a way John's toe *appears* to John without making any claim on the condition of his toe. John may come to believe or judge that something is (physically) wrong with his toe as expected. But this is not a proper perceptual judgment directly based on the pain experience — it's inferential but usually quite habitual. It's very different than the genuine perceptual belief [this is red] directly formed on the basis of having a visual experience of a red apple under normal circumstances.²³ *P* is a phenomenal way of appearing — it presents body parts “painfully” without attributing a quality (objective or otherwise) to those body parts.

Our discussion of John's pain in his toe and his utterance has revealed the remarkable peculiarities of L-utterances and L-judgments. Pain scientists and clinicians like everybody else routinely use this language without suffering from the apparent conflict it seems to present in light of the IASP definition of pain as a subjective experience. (6) gives us the truth-conditions of L-utterances by making essential use of the phenomenal character of pains as subjective experiences — as the IASP definition presupposes.²⁴ If we take the bodily conception of pain to be denoting the pains reported in L-utterances, i.e., pains located in body parts, then the bodily conception of pain turns out to be perfectly compatible with the experiential conception of pain as stressed by the IASP definition. Indeed, they are not only compatible, but (6) requires the experiential conception.

As with options in (5), it is possible to express (6) in terms of the language of inverse intentional properties. With his true utterance, John locates a bodily pain in his toe. Under (6), the “property” located is the inverse property of [the toe's] appearing a certain way to John (“appearing painfully to”). Under the naturalistic implementational framework I prefer, this amounts to his toe's instantiating the inverse property of being the referential target of John's pain experience with respect to which a proprietary predicate (implementing the same experience) is being activated in John's nociceptive system. We noted before that in having a pain in his toe, John is confronted with an appearance-

²² Except that (6) is meant to be compatible with non-primitivism — see fn. 15.

²³ Judgments [this is damaged] and [this hurts] made consequent upon having a pain experience where ‘this’ refers to a body part are very different judgments with radically different truth-conditions.

²⁴ See my (2017c) for an in-depth discussion of the main driving force behind the IASP definition of pain.

dependent quality of his toe, and indeed he himself would take this quality to be appearance-dependent (at least after reflection) and to be metaphysically independent of any physical condition of his toe.²⁵ (6) tells us how this can be so — making the bodily conception harmonious with the experiential conception of pain as embodied in the IASP definition.

How does (6) compare to (4)? There is a deflationary way of reading (4) so that it becomes compatible with (6). (4) quantifies over a subjective quality in the toe outside of the scope of the intentional verb ‘sense’ (or, ‘is aware of’), but then makes the instantiation of the quality metaphysically dependent on its being sensed. This generates an uneasy circularity. Whether the circularity is vicious depends on further claims one might want to make about the situation. One way to understand (4) is to identify the *act* of sensing with the instantiation of the pain quale *P* by the pain experience: the instantiation of *P* in John’s experience (in relation to his toe) *is* the act of his sensing the toe painfully. The act of sensing his toe painfully is then to be identified as a whole with the instantiation of a subjective quality in John’s toe. If we do this and tolerate the *loose* talk of John’s sensing this quality, then (4) and (6) become compatible indeed. But under this interpretation, it *is* loose talk to speak of sensing or being aware *of* a subjective quality in one’s toe: strictly speaking, John is not sensing any quality in the sense in which we use this phrase to talk about sensing sensible properties of extra-mental objects or events such as sensing the colours or vibration of objects or surfaces. Of course, when John reflects on its experience,

²⁵ John is representative of the regular folk. Even if we put the folk aside, there is no doubt that pain scientists and clinicians think of bodily pains along these lines. Adam Shriver and I have been running various pilot studies among pain scientists and clinicians surveying them about their linguistic intuitions. The results so far quite robustly support the claim I’ve just made. They overwhelmingly choose options (4–6) as to what makes John’s utterance true. We’ll soon be preparing the results of our experimental work for publication.

This is probably as good a place as any to say a few words on the experimental work done by Kevin Reuter and Justin Sytsma and his colleagues surveying the folk about pain attributions (e.g., Reuter et al. 2014, Sytsma & Reuter 2018, among other works). First, I am quite familiar with this body of work. Second, I am not at all convinced that it shows something philosophically significant. At best, this body of work continues to buttress the observation (made by a lot of people) that the folk conception of pain involves confusions and conceptual tensions. The authors’ claim that the results show that the folk conception of pain heavily prefers the bodily conception is not persuasive given two facts. One is that everybody agrees that the folk conception of pain has these two (bodily and experiential) strands in it (see my 2005/2013 for explicit statement and elaboration). The other is that depending on how you set up the vignettes and questions, you can get results supporting either strand — this is not surprising at all. To see a detailed and pretty convincing criticism of their work along these lines, see Borg et al. (forthcoming). Furthermore, it is not at all clear what Sytsma et al. think the bodily conception of pain is that they attribute to Folk. They *seem* to think that whatever it is, it is incompatible with the experiential conception of pain. (They also misattribute to me the view that these two strands are incompatible.) Even more vexingly, they seem to think that the bodily conception, whatever it is, is not captured by anything along the lines of (2)–(3) since they seem aware that the referred pain cases are cases where the sincere bodily pain reports are regularly taken to be true. But, then, it is not quite clear why they think that these two conceptions are incompatible with each other. If they are not incompatible, it is not clear what they take their work to show philosophically. So, here and elsewhere, I set their work aside. I argue in the present work that these two conceptions are not only compatible with each other but also the bodily conception requires the experiential conception. Furthermore, as mentioned above, we have robust experimental results to show that the pain scientists and clinicians’ dominant conception is experiential — just as the IASP definition says.

it will seem to him *that* his toe has a subjective quality instantiated in it. But this is *doxastic seeming* immediately prompted by his experience. There is no other way for John to reflect on his own experience. But this is perfectly in line with the truth of his *judgment* (and L-utterance) when he locates pain in his toe just as (6) says: he is sensing his toe painfully — thus there is a subjective quality (a pain) instantiated in his toe, making his judgment or L-utterance true. Furthermore, John’s judgment is appropriately based on his so sensing the toe.

Before concluding this work, I would like to briefly discuss a peculiar attempt to save representationalism. Given the implementational framework I have in mind for (6), one might be tempted to say the following.²⁶ The proprietary predicate activated (with respect to John’s toe) in John’s nociceptive system *represents* the toe as being the (semantic) target of its own activation, which we identified with the toe’s having a pain in it. Let ‘*#P(t)*’ denote the *activation* of this sensory predicate targeting the toe referred by the sensory referential device *t*. Then the suggestion becomes: *#P(t)* is *true* IFF *t* successfully refers to John’s toe and is the filler of the argument-place in *#P()*. The truth of the right-hand side is what makes it the case that *#P()* *represents* itself as being instantiated with reference to John’s toe (thus, making *#P(t)* *true*): this in turn is what makes it the case that John has a pain in his toe. So, it seems like, we can have our representationalist cake and eat it too if we tolerate self-representational and self-verifying sensory predicates.

I have various reasons for wanting to avoid taking this line. Three of them stand out as worthy of mentioning here quickly. First, I am very fond of an information-theoretic psychosemantics à la Dretske for sensory and perceptual mental predicates (see Aydede & Güzeldere 2005). According to information-theoretic principles, no event can carry information about its own occurrence.

Second, and relatedly, for a psychosemantic naturalist like me, a self-verifying mental predicate is not only mysterious but is also an anathema to taking intentionality seriously enough to attempt to naturalize it: if you take the notion of (attributive) representation seriously, it ought to be the notion of something that can be misapplied, misattributed — a representation that cannot misrepresent is no representation at all.

Third, even if we put these worries aside, taking this line forgets what the primary impetus was behind representationalism about sensory experiences in the first place: namely, to avoid quantifying over “irreducibly psychic” items like sense-data or Qualia. But a representationalist who claims that a pain in the toe is nothing but a self-representing and self-verifying mental quality instantiated in one’s experience semantically targeting the toe has to advert to a non-representational quality of the experience (*qua* sensory *predicate* — not content — as should be clear from the statement of the bi-conditional above), which is anathema to any transparency claims that representationalists are fond of making about sensory experiences (see my 2009, 2017b). Moreover, any representationalist who likes to take this line with the intention of putting all bodily sensations on the same naturalist footing would make a mockery out of their claim that in a sensory experience like pain in

²⁶ Many thanks to Mike Martin who has pushed me to take this line of thinking seriously.

the toe, John is detecting, sensing, feeling a condition of his toe, in the same way in which John would be detecting or sensing the vibration in his toe upon, say, starting the engine of his sailboat. On the other hand, there is a way of taking this representationalist proposal as representationalist only *by courtesy*: all that happens is that the sensory predicate $\#P()$ is activated with t as the filler of its argument-place referring to the toe, and *then* it's claimed that $\#P()$ represents the toe as being the target of its own activation. This sounds like a claim tacked at the last minute in virtue of one's representationalist predilections, not quite clear what work it's supposed to be doing. We don't need to make this further claim to make sense of (6). It complicates things unnecessarily and has unwelcome consequences: if the implementational framework is generalized to all sensory experiences, it makes *all* sensory experiences have a layer of self-representational and self-verifying content that "mentions" non-intentional ("syntactic") features of experiences that are epistemically available to sensing, perceiving, and introspecting agents of a certain complexity. This should be alarming to direct realists and (ironically) to naturalist representationalists themselves — see Aydede (2017b).

Finally, referred pains pose problems for many proposals discussed above. In particular, (2)–(3) and the Shoemaker's version of (5iii) are unable to explain the truth of sincere L-utterances ascribing referred pains to otherwise healthy and functioning body parts. Other versions of (5) can accommodate referred pain ascriptions but they make pain experiences semantically mistaken — even though, as we've seen, these mistakes do not at the end affect the assignments of correct truth-conditions to L-utterances reporting referred pains. (6) has none of these problems. If John's pain in his toe is a referred pain, then John's sincere utterance is true and there is no semantic mistake in John's experience. The referential term in John's experience, t (perhaps realized in the "toe" part of the somatosensory homunculus in John's primary somatosensory cortex), successfully refers to his toe: it presents *his toe* to John. The somaesthetic informational channel necessary to referentially connect his toe to John's cortex is intact, and there is for the most part good information flowing through this channel despite the glitch in John's lower spine. This glitch causes the referential device, t , to activate a neural pattern, $\#P$ (realizing the pain quale P), which makes John's toe appear to him *painfully*. The experiential reference to the toe is successful, but it's unfortunate that the experience presents this toe in the *way* it does: this presentational *way* (the instantiation of P) in this particular case doesn't carry any information about the physical condition of the successfully presented toe. So John's experience is misleading but not mistaken.²⁷

²⁷ Phantom limb pains are pains felt, as usually described, "in missing body parts." Suppose John is missing his toe but still feels a pain in his "phantom toe." Then, John doesn't have a *pain in his toe*. If John didn't know he was missing his toe and sincerely uttered that he had a pain in his toe, then he would have said something false. Here we have a reference failure in John's experience: the informational channel is completely broken. But clearly it doesn't follow that John doesn't experience pain. He still is in genuine pain in so far as he has an experience that makes it seem to him that his toe is presented to him *painfully*. The IASP definition and (6) agree that pains are subjective experiences — however the syntactic form of such experiences structures the phenomenology.

4 Concluding Remarks

There are various other virtues of (6). Unlike representationalist or perceptualist accounts, it frees pain experiences from committing mistakes or error. Unlike the projectivist options in (5ii) and (5iii), there are no representational mistakes anywhere in John's pain system according to (6) and the implementational framework I prefer. One robust mark of projectivism, as we've seen, is that it makes experiences under consideration massively illusory or somehow mistaken. Nothing of this sort happens with (6). So (6) is not projectivist. Indeed (6) comports well with the traditional understanding of pain as not representational. In all these respects, the account of pain attributions to body parts both in language and thought provided by (6) is superior to its competitors. It solves the puzzle with which we started this essay — and it does that without attributing deep confusion or incoherence to the folk or to pain scientists or clinicians. There is no real paradox in the ordinary, clinical, or scientific understanding of pain — just a conceptual tension of the sort a careful philosophical analysis would clarify and relieve.

The importance of pain experiences as in other intransitive bodily sensations is that their peculiarities tell us something very important about the metaphysics of the conscious perceptual experience in general, something that the contemporary philosophy of mind, with its heavy reliance on representationalism and visual perception, has lost sight on.²⁸

References

Aydede, M. (2006). The Main Difficulty with Pain. In M. Aydede (Ed.), *Pain: New Essays on Its Nature and the Methodology of Its Study*. Cambridge, Massachusetts: MIT Press, 123–136.

Aydede, M. (2009). Is Feeling Pain the Perception of Something? *Journal of Philosophy*, 106(10), 531–67.

Aydede, Murat. (2005/2013). Pain. *The Stanford Encyclopedia of Philosophy* (Spring 2013 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2013/entries/pain/>

Aydede, Murat (2014). How to Unify Theories of Sensory Pleasure: An Adverbialist Proposal. *Review of Philosophy and Psychology*, 5(1), 119–33.

²⁸ The implications of this sort of view for a general understanding of conscious perceptual experience and its introspection are explored in other work — see Aydede & Güzeldere (2005) and Aydede (2017b). Different versions or portions of this work were delivered at the University of Oxford and the University of Reading on separate occasions. I would like to thank the audiences. In particular, many thanks to Mike Martin and Emma Borg for their helpful comments and questions on these occasions. Thanks also to Adam Bradley, Laurenz Casser, Jonathan Cohen, Linda Davies, Matt Fulkerson, Joshua Gert, Adam Shriver, and Andrew Wright for their insightful comments on earlier versions of this work.

Aydede, Murat (2017a). Pain: Perception or Introspection. In *The Routledge Encyclopedia of Philosophy of Pain*, edited by J. Corns (Routledge).

Aydede, Murat (2017b). Is the Experience of Pain Transparent? Introspecting Phenomenal Qualities. *Synthese*. <doi: 10.1007/s11229-017-1528-3>

Aydede, Murat. (2017c). Defending the IASP Definition of Pain. *The Monist* 100 (4), 439–64.

Aydede, Murat (2017d). Critical Notice of Colin Klein's *What the Body Commands: The Imperative Theory of Pain* (MIT 2015). Presentation for the *Author-Meets-Critics* session on Klein's book at the Eastern APA Division Meeting, Baltimore, January 4, 2017. (Accessible at <<https://philpapers.org/rec/MURCNO>>).

Aydede, Murat (2018). A Contemporary Account of Sensory Pleasure. In *Pleasure: A History*, edited by Lisa Shapiro. New York: Oxford University Press, 239–266.

Aydede, Murat, and Fulkerson, M. (2014). Affect: Representationalists' Headache. *Philosophical Studies*, 170(2), 175–198.

Aydede, Murat, and Güven Güzeldere. (2005). Cognitive Architecture, Concepts, and Introspection: An Information-Theoretic Solution to the Problem of Phenomenal Consciousness. *Noûs* 39(2), 197–255.

Bain, D. (2003). Intentionalism and Pain. *Philosophical Quarterly*, 53(213), 502–522.

Bain, D. (2007). The Location of Pains. *Philosophical Papers*, 36(2), 171–205.

Bain, D. (2013). What Makes Pains Unpleasant? *Philosophical Studies* 166 (1 Supp), 69–89.

Borg, Emma, Richard Harrison, James Stazicker, and Tim Salomons (forthcoming). Is the folk concept of pain polyeidic? *Mind and Language*.

Bradley, Adam (2018.) The Paradox of Pain. Chapter in PhD Dissertation, UC Berkeley.

Clark, Austen (2000). *A Theory of Sentience*. Oxford University Press.

Chalmers, David J. (2006). Perception and the fall from Eden. In Tamar S. Gendler & John Hawthorne (eds.), *Perceptual Experience*. Oxford University Press, 49–125.

Hill, C. (2006). Ow! The Paradox of Pain. In M. Aydede (Ed.), *Pain: New Essays on Its Nature and the Methodology of Its Study*. Cambridge, Massachusetts: MIT Press, 75–98.

Hill, C. (2009). *Consciousness*. Cambridge, UK: Cambridge University Press.

- Hyman, John. (2003). Pains and Places. *Philosophy* 78 (01), 5–24.
- IASP. 1979. IASP Pain Terms — 1979. *Pain*, 6: 247–252. (2011 on-line version: <<http://www.iasp-pain.org/Education/Content.aspx?ItemNumber=1698>>)
- Klein, C. 2015. *What the Body Commands: The Imperative Theory of Pain*. Cambridge, Massachusetts: MIT Press.
- Merskey, H. (2008). History and Definition of Pain. In *Chronic Pain: a Health Policy Perspective*, edited by S. Rashiq, D. Schopflocher, P. Taenzer and E. Jonsson (WILEY-VCH Verlag GmbH & Co. KGaA, Weinheim, 1–10).
- Pautz, Adam. (2010). Do Theories of Consciousness Rest on a Mistake? *Philosophical Issues* 20: 333–67.
- Reuter, K, Phillips, D., and Sytsma, J. (2014): Hallucinating Pain. In *Advances in Experimental Philosophy of Mind*, ed. Justin Sytsma, London: Bloomsbury.
- Shoemaker, S. (1994). Self-Knowledge and ‘Inner Sense’: Lecture III: the Phenomenal Character of Experience. *Philosophy and Phenomenological Research*, 54(2), 291–314.
- Shoemaker, S. (2000). Introspection and Phenomenal Character. *Philosophical Topics*, 28(2), 247–273.
- Sytsma, J., and Reuter, K. (2018). Unfelt pains. *Synthese*. Online First.
- Tye, M. (2006a). Another Look at Representationalism about Pain. In M. Aydede (Ed.), *Pain: New Essays on Its Nature and the Methodology of Its Study*. Cambridge, Massachusetts: MIT Press, 99–120.