

MIT Open Access Articles

Hmm... Hill on the paradox of pain

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Byrne, Alex. "Hmm... Hill on the Paradox of Pain." *Philosophical Studies* (2011): p. 1-8.

As Published: <http://dx.doi.org/10.1007/s11098-011-9811-5>

Publisher: Springer Science + Business Media B.V.

Persistent URL: <http://hdl.handle.net/1721.1/70542>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



Hmm... Hill on the paradox of pain

Alex Byrne

1. Pain perception

Bodily sensations, for instance pains, are often thought to pose insuperable difficulties for representational theories of consciousness. Chris is never one to dodge a problem, and chapter 6 of his splendid and instructive *Consciousness* is devoted to outlining a perceptual/representational theory of pain, in the tradition of Armstrong and Pitcher.

Perceptual theories of pain are not the same as representational theories of pain. According to a perceptual theory of pain, when one has a pain in one's toe (for example), one perceives some sort of disturbance in one's toe. (At least in a typical case: the perceptual theorist will want to allow for illusions and hallucinations of such disturbances.¹) Pain perception is *interoception*—perception specialized for delivering information about one's own body, like proprioception and the vestibular sense.

According to a representational theory of pain, the felt quality of pain, or the “phenomenal character” of pain, can be fully explained in representational or intentional terms. A perceptual theory does not entail a representational one, anymore than a perceptual theory of vision (which of course everyone holds) entails a representational one (which not everyone holds). Neither does the converse entailment hold. One may hold a representational theory of imagery, say, without thereby turning imagery into perception. I will set aside the issue of representationalism, and concentrate here on what Chris calls the “paradox of pain”, which he presents chiefly as a problem for perceptual theories of pain.

Perception involves dedicated mechanisms of sensory transduction that result in the delivery of ecologically useful information about the perceiver's environment. Vision, to take the most prominent example, delivers a large variety of information about the

¹ A perceptual theorist might even hold that the experience of pain is *never* veridical (as some hold that experiences of color are never veridical). This “eliminativist” variant of the perceptual theory will be ignored here.

shape, motion, texture, and color of the scene before the eyes. The extra-bodily environment is not the only part of the environment of interest—the state of the perceiver herself is also important. A perceiver can turn her external senses towards herself: you can check the position of your feet by looking, or palpate a lump on your head, for instance. But other more direct methods are useful, and so it is no surprise that there are specialized mechanisms for this purpose: proprioception informs us about the position of our limbs, and the vestibular sense informs us about our balance. Such information about our bodily states is crucial for successful action, of course.

Similarly, it's no surprise that a car has systems designed to detect its internal states (the gas gauge and tachometer, for example)—they are useful for much the same reason. Further, the gas gauge (measuring the internal environment), and the outside air temperature sensor (measuring the external environment) deliver their proprietary messages in the same calm unobtrusive manner, not making much of a song and dance about it.

Sometimes conditions threatening the car's well-being arise that demand quick action—near-empty tank, low oil pressure, etc. Cars have systems to detect those conditions as well, and it's no surprise that the message delivery is more strident and eye-catching—the dashboard warning light. (There is even the maximally vague “check engine” light, apt to induce disturbing sense of unease. *Something* is wrong, but what?)

Pursuing the parallel, one might expect organisms to have their own versions of dashboard warning lights, and there are no better candidates than pains. And since a dashboard warning light is not just a pretty glow, but a messenger, like the speedometer, these armchair biological reflections motivate a perceptual view of pain.

Consciousness contains many persuasive considerations (some drawn from the empirical study of pain) in favor of a perceptual view, so in what follows I will take a perceptual view for granted.²

Some terminology before we get going. Let *painful disturbances* be the sorts of items that (according to the perceptual theorist) one perceives via one's pain-perception system—in the case of having a pain in one's toe, the disturbance that is (or at any rate

² For important dissent subsequent to the publication of *Consciousness*, see Aydede 2009.

appears to be) in the toe. (For simplicity I shall ignore cases where the perceived object is a body part, as in ‘My toe is painful/hurts’.) Whether such disturbances are to be identified with anything physical is not our present focus.³ *Experiences of pain* are the states of being ostensibly aware of such disturbances.

A bare perceptual view of pain is not committed to any view about the semantics of ‘pain’ and its cognates. One could hold a perceptual view and deny that we *ever* use locutions involving the word ‘pain’ to refer to painful disturbances. That sounds odd, but it is actually Chris’s final position. ‘Pain’, Chris thinks, does *not* refer: he “[i]n effect...denies that there is such a thing as pain” (Hill 2009: 224). A fortiori, ‘pain’ does not refer to painful disturbances.

Still, on the assumption that the word ‘pain’ does refer, it would be quite peculiar if painful disturbances were perceived and the folk never used the word ‘pain’ to pick them out. We usually have commonplace words and phrases that can be readily used to describe the objects of perception (‘yellow banana’, ‘C-sharp’, ‘rough surface’, etc.), so why not here? (If painful disturbances turn out to be, say, certain kinds of bodily insults, the theorist might refer to them in those terms. But the Plain Man in pain is not in a position to do that.) So let *the simple perceptual view of pain* be the view that painful disturbances are perceived and that certain locutions involving ‘pain’ are used to refer to them.

2. The paradox of pain

The paradox is generated by the conjunction of the simple perceptual view of pain and either one of the following two principles:

(A) If x is in pain, then it seems to x that x is in pain, in the sense that x has an experiential ground for judging that x is in pain

(B) If it seems to x that x is in pain, in the sense that x has an experiential ground for judging that x is in pain, then x really is in pain (169).

³ According to Chris, painful disturbances are “bodily disturbances”, “peripheral disturbances involving actual or potential damage” (Hill 2009: 177). That is not being assumed here.

Successful perception in general requires the cooperation of the object perceived and the perceptual mechanism. On the one hand, the object might not affect the mechanism, or might not affect it in the right way. The mere presence of a red tomato, or the mere crossing of my legs, is not sufficient for the perception of the tomato, or my crossed legs. On the other hand, the mechanism might be affected in the way it is in cases of perception, but not by the object. So if pains are perceived (as the simple view has it), then both cases should be possible: unperceived pain, and pain hallucinations. I think Chris would accept the following gloss on (A) and (B), respectively: all pains are perceived, and there are no pain hallucinations. Thus Chris claims that (A) and (B) are “two folk principles that preclude an appearance/reality distinction for pain” (186), ruling out the simple perceptual view. That is the paradox of pain.

The simple perceptual view of pain is plausible. But are (A) or (B) plausible too? According to Chris, plausibility is something of an understatement:

[W]e regard it as *absurd* to say that an agent is in pain in circumstances in which the agent is not aware of the pain, and...we regard it as *absurd* to say that an agent is not in pain in circumstances in which it seems to the agent that he or she is in pain. This suggests that we think of (A) and (B) as necessary truths—that is, as holding either because of deep metaphysical facts about our awareness of pain, or because of the a priori structure of our concept of pain. (170)

Chris does go on to say that there is some recalcitrant evidence—we think that pains can wake us up, for example—but dismisses it as not “represent[ing] a dominant strand in our commonsense conception of pain” (171).

As I said, I agree with Chris that painful disturbances are perceived. I also agree that (A) and (B) are plausible (in fact, I think they are true). I don’t, though, think this generates a paradox for the simple perceptual view. There is a closely related puzzle in the vicinity, but as far as I can see this has nothing in particular to do with pain.

3. The status of (A) and (B)

According to Chris, the principles (A) and (B) “are partially constitutive of the concept of pain” (186). I think that Chris understands this to have the following consequence: anyone possessing the concept of pain believes (A) and (B), and moreover believes them

“to enjoy a kind of necessity” (171)—or at any rate is disposed to have these beliefs on contemplating these two principles. If that is right, then by his own lights Chris cannot respond to the paradox simply by denying (A) and (B), since he certainly possesses the concept of pain. And indeed he does not. He thinks that the paradox shows that the concept of pain is “semantically incoherent” (190), somewhat as the concept of a round square is semantically incoherent. A round square must be round, and also must be square, but nothing can be both. Similarly, according to Chris, “Folk psychology...requires us to think of pains as objects of experiential awareness” (188), which makes an appearance/reality distinction all but inevitable; but folk psychology also insists that there is no such distinction. Chris’s positive recommendation is that we set “[the concept of pain] aside” and replace “it with three new concepts that have separate conceptual roles” (190), specifically concepts for painful disturbances, experiences of pain, and pain affect. Although Chris doesn’t explicitly deny (A) and (B), he does speak of “rejecting” them (188)—if this doesn’t amount to denying (A) and (B), what does it mean? One might have expected Chris to *endorse* (A) and (B), while emphasizing that they don’t imply that anyone really *is* in pain.

Chris argues for the conclusion that the principles (A) and (B) “are partially constitutive of the concept of pain” in two steps. First, he argues that *if* “the folk concept of pain is governed by” a certain pair of rules, then (A) and (B) are partially constitutive of the concept of pain. Second, he argues that the folk concept of pain *is* governed by these two rules, on the grounds that “we have a very powerful motive for having a concept” that is so governed.

Chris states the two rules as follows:

(R1) One is fully and unqualifiedly entitled to form a first person judgment of the form *I am in pain* if one is currently having an experience of pain.

(R2) One is not entitled to form a first person judgment of the form *I am in pain* unless one is currently having an experience of pain. (183)

One question that arises at this point, and which is not explicitly addressed in the book, is just what Chris means by saying that a concept is “governed by” such-and-such rules.

Moreover, whatever the answer to this, (R1) and (R2) are not in fact cast in the form of rules ('Do so-and-so', 'You should do so-and-so', etc.).

However, Chris's intent is not hard to discern, and for present purposes a close enough approximation to it is, I think, as follows. (See Hill 2005: 82-3.) Our concept *pain* is governed by (R1) and (R2) just in case, respectively:

(G1) Someone possessing the concept *pain* is disposed to form the first person judgment *I am in pain* when she is currently having an experience of pain, no matter what else she believes at the time.

(G2) Someone possessing the concept *pain* will not form the first person judgment *I am in pain* unless she is currently having an experience of pain, no matter what else she believes at the time.

According to Chris, if "a subject is conforming to [(R1) and (R2)], then he or she will have absolutely no motivation for distinguishing between the appearance of pain and the reality of pain" (184). That is, if (G1) and (G2) are true, then anyone who possess the concept *pain* "will have absolutely no motivation for distinguishing between the appearance of pain and the reality of pain". As Chris puts it a little later, the truth of (G1) and (G2) explains "why it is not apparent to us from the perspective of folk psychology that it is possible to draw an appearance/reality distinction with respect to pain" (185). This is Chris's way of restating the claim that (G1) and (G2) show that (A) and (B) are "partially constitutive of our concept of pain". So do they?

Take (G1). It is supposed to show, specifically, that (B) is "partially constitutive of our concept of pain", which I am taking Chris to mean (roughly) that anyone who possesses the concept *pain* will believe it to be necessary that whenever there is the appearance of pain, there is pain.

I possess the concept *pain*, so if (G1) is true, I am disposed to judge that I am in pain whenever I am having an experience of pain, no matter what else I believe. One immediate problem is that the mere *truth* of (G1) is surely not going to suggest any (B)-like principle to me—I also need to *know* (or believe) that (G1) is true. So let us further suppose (since this can only help Chris) that I know that I have this disposition. It is not clear to me why this implies that I will believe the appearance of pain is sufficient for the

reality of pain. First, I sometimes think I have belief-forming dispositions that are not entirely reasonable. Perhaps this is a case in point: I am just overly gullible where the appearance of pain is concerned. Second, what about the appearance of pain in others? Why could I not regard myself as one of the privileged few for whom the appearance of pain is an infallible guide? These seem live options, at least—why must I dismiss them out of hand? Third, where is the *necessity* coming from? Even if the argument shows that I believe that whenever there is the appearance of pain, there is pain, why must I regard the negation of this as absurd? Why can't I think of it as contingent?

Similar remarks go for Chris's attempt to argue from (G2) to the conclusion that (A) is partially constitutive of the concept of pain.

So I am not convinced by the first step of Chris's argument, that *if* the folk concept of pain is governed by the two rules (R1) and (R2), then the principles (A) and (B), as Chris understands them, are partially constitutive of the concept of pain. What about the second step, that the folk concept of pain *is* governed by these two rules? Chris writes:

Is it true that the folk concept of pain is governed by (R1) and (R2)? I think we can see that it is true, and also understand why it is true, by noticing that we have a very powerful motive for having a concept that can be used to keep track of the experience of pain—that is, to encode and store information about the experience. The motivation comes from the causal powers that are possessed by the experiences of pain, or, equivalently, by the somatosensory representations that constitute such experiences. It is these somatosensory representations that directly control the activity in the anterior cingulate cortex and in the limbic system that determines our affective response to pain...And of course, it is pain affect that makes pain so important to us. Without the response, pains would be of no more interest to us than tingles and mild sensations of pressure. (185)

We do, Chris acknowledges, have a motive for having a concept of painful disturbances—he identifies them with what he calls “bodily disturbances”⁴—but:

⁴ See note 3 above.

Bodily disturbances have the power to cause pain affect only insofar as they have the power to cause somatosensory representations. Accordingly, it is much less important that we have a concept that can be used to keep close tabs on bodily disturbances than that we have a concept that can be used to keep close tabs on representations of the relevant sort. (185)

Chris is quite explicit about the upshot: the folk concept of pain (or, more simply, the concept *pain*) is governed by (R1) and (R2). We can break this down into two parts: (i) we have a concept C that is governed by rules that are structurally like (R1) and (R2), and (ii) C = the concept *pain*.

What Chris appears to offer in support of (i) and (ii) is an explanation of why we have “a concept that can be used to keep track of the experience of pain”. In brief, we have that such a concept because “we have a very powerful motive for having [it]”. Suppose that explanation is right and let the concept that keeps track of the experience of pain be C*. What’s the argument that C* is governed by rules like (R1) and (R2), or that C* is the concept *pain*?

Concerning the second question, the argument cannot be that the concept *pain* is the *only* candidate to be C*, because one concept we have “that can be used to keep track of the experience of pain” is, of course, the concept *experience of pain*. And by Chris’s own lights, the concept *experience of pain* really is *another* candidate for C*—that concept is *not* the concept *pain*, because the concept *experience of pain* has application, unlike the concept *pain*.

Chris’s discussion here is something of a double-edged sword. As he points out, we have an interest in keeping tabs on experiences of pain, and *also* on painful disturbances. (Remember the dashboard analogy. When the check engine light goes on, we want it to go off, because it’s a big distraction. But the fault in the engine is also of great importance.) Presumably we also have an interest in keeping tabs on experiences of pain and painful disturbances without falling into semantic incoherence. So why don’t we?

4. (A) and (B) revisited, and the puzzle of pain

Let us go back to Chris's two principles that are supposed to generate the paradox of pain.

(A) If x is in pain, then it seems to x that x is in pain, in the sense that x has an experiential ground for judging that x is in pain

(B) If it seems to x that x is in pain, in the sense that x has an experiential ground for judging that x is in pain, then x really is in pain

First, I shall argue that both (A) and (B) are true—or at any rate true when read in a very natural way. Then I shall argue that there is no conflict with the simple perceptual view. Chris's paradox of pain doesn't completely evaporate, though; it leaves a residue that I will call the *puzzle of pain*, and I shall end with a brief discussion of that.

When we say that x *feels pain*, or is *in pain*, what are we saying? If x 's phantom limb is playing up again, surely he feels pain. That strongly suggests that to *feel* pain, or (what seems to be equivalent) to be *in pain*, is simply to enjoy an experience of pain. 'Feels hot' provides a close parallel, as Armstrong pointed out (1968: 315). If one *feels* hot, then one does not have to *be* hot. Similarly if one's hand feels hot. For one's hand to feel hot is simply for one to enjoy an experience of heat-in-the-hand.

Now what does 'x has an experiential ground for judging that x is in pain' mean, in the statement of (A) and (B)? Clearly it means that x has an experience of pain. And, as just argued, to have an experience of pain is to be *in pain*. So, since 'it seems to x that x is in pain' is supposed to be understood as 'x has an experiential ground for judging that x is in pain', (A) and (B) are both equivalent to the tautologous:

(☺) If x is in pain, then x is in pain.

Now from the fact that there are locutions involving 'hot' that refer to experiences of heat, it does not follow that there are no locutions involving 'hot' that refer to actual thermal phenomena. Likewise, from the fact that there are locutions involving 'pain' that refer to *experiences of pain*, it does not follow that there are no locutions involving 'pain' that refer to *painful disturbances*. And (as Chris himself points out) such locutions are

easy to find, namely those locating pains in a part of the body—there is a pain in my foot/I have a pain in my foot, etc.

So as far as I can see, we can have our cake and eat it too. (A) and (B) are true, and so is the simple perceptual view of pain: painful disturbances are both perceived and talked about.

Finally, the residual “puzzle of pain”. Earlier, I said that Chris would accept the following gloss on (A) and (B): all pains are perceived, and there are no hallucinations of pain. I have argued, in effect, that that this isn’t the right gloss on (A) and (B). But the gloss *itself* seems right. If x has a pain in his foot, he must feel, or be aware of, the pain in his foot. And if x is ostensibly aware of a pain in his foot, there may be no foot, but there must be a pain that x is aware of. Or so we are inclined to think.

According to the simple perceptual theory, these claims we are inclined to accept are false; at the very least, they are not necessarily true, or “constitutive of our concept of pain”. Fortunately for the simple perceptual theory, Chris has not argued that they are.

A simple perceptual theorist needs to explain why we have these mistaken inclinations.⁵ But whatever their source, I doubt it has anything to do with the special features of pain. We are equally apt to collapse appearance and reality for afterimages and non-painful tingles, for example. As anyone who has taught intro philosophy of mind knows, many undergraduates find the view that afterimage-experiences are hallucinations very hard to accept.

And finally, Chris faces his own version of the puzzle of pain. For him, it is a puzzle expressed in our jargon of ‘painful disturbances’, rather than the ordinary word ‘pain’, but it is no less puzzling for that. Consider the following two claims. If x has a painful disturbance in his foot, he must feel, or be aware of, the painful disturbance in his foot. And if x is ostensibly aware of a painful disturbance in his foot, there may be no foot, but there must be a painful disturbance that x is aware of. Suppose Chris is right about the incoherence of ‘pain’. Still, don’t those two claims strike us as *true*?

Chris’s penetrating chapter on pain brings us good and bad news: a perceptual view of pain is true, but our ordinary thought and talk about pain is fundamentally

⁵ For some helpful suggestions see Ganson and Ganson 2010.

defective. I have tried to soften the bad news a bit: we are confused about pain, but ordinary language is perfectly in order as it is.⁶

⁶ Thanks to Murat Aydede for comments.

Bibliography

Armstrong, D. M. 1968. *A Materialist Theory of the Mind*: Routledge and Kegan Paul.

Aydede, M. 2009. Is feeling pain the perception of something? *Journal Of Philosophy* 106: 531-67.

Ganson, T., and D. Ganson. 2010. Everyday thinking about bodily sensations. *Australasian Journal of Philosophy* 88: 523-34.

Hill, C. 2005. Ow! The paradox of pain. *Pain*, ed. M. Aydede. Cambridge, MA: MIT Press.

— — —. 2009. *Consciousness*. Cambridge: Cambridge University Press.