

# The Main Difficulty with Pain: Commentary on Tye

Murat Aydede  
University of Florida  
Philosophy Department  
May 2004

Consider the following two sentences:

- (1) I see a dark discoloration in the back of my hand.
- (2) I feel a jabbing pain in the back of my hand.

They seem to have the same surface grammar, and thus *prima facie* invite the same kind of semantic treatment. Even though a reading of ‘see’ in (1) where the verb is not treated as a success verb is not out of the question, it is not the ordinary and natural reading. Note that if I am hallucinating a dark discoloration in the back of my hand, then (1) is simply false. For (1) to be true, therefore, I have to stand in the seeing relation to a dark discoloration in the back of my hand, i.e., to a certain surface region in the back of my hand marked by a darker shade of the usual color of my skin, a certain region that can be seen by others possibly in the same way in which I see it. Also note that although the truth of (1) doesn’t require the possession of any concept by me expressed by the words making up the sentence, my uttering of (1) to make a report typically does — if we take such utterances as expressions of one’s thoughts. So my seeing would typically induce me to identify something in the back of my hand as a dark discoloration. This is a typical case of categorization of something under a concept induced by perception. Of course, my uttering of (1) does more than attributing a physical property to a bodily region, it also reports that I am seeing it.

What has to be the case for (2) to be true? Like (1), (2) invites us to follow its surface grammar and treat it as expressing a perceptual relation between me and something else which has a bodily location, viz., a jabbing pain in the back of my hand. The difficulties with taking this route are familiar, so I won’t repeat them here except to say that whatever the true analysis of sentences like (2) turns out to be, one thing is clear: the truth-conditions of (2) put no constraints whatsoever on how things physically are with my hand. Anyone who has a sufficient mastery of our ordinary concept of pain has no difficulty whatsoever in understanding how (2) could still be true even though there is nothing physically wrong with my hand. So if the truth of (2) is taken to imply attributing a property to the back of my hand, this property is not a physical property. For many with naturalistic leanings like me, this is one of the main reasons for not taking (2) as primarily making a property attribution to a bodily region. So when I utter (2) to make a report and appear to identify something in the back of my hand as a jabbing pain, whatever else I am doing, I am not attributing a property or a condition (physical or otherwise) to the back of my hand and saying that I am feeling it.<sup>1</sup>

---

<sup>1</sup> There are different ways of reporting pain in body regions. Many involve subtleties for conveying contextual information. But insofar they are genuine reports of pain in body parts following our ordinary

If this is not what I am doing, what is it that I am doing? What is the proper analysis of sentences such as (2)? On the basis of considerations harnessed so far, one is tempted to argue in the following way. Every genuine case of perception allows reporting an instance of perception in the relevant modality by sentences similar to (1), where the perceptual verb is used dominantly as a success verb. This is for good reason: perception is essentially an activity whereby one gathers information about one's environment in real time (including one's internal bodily environment of course). So it is not surprising that the dominant form of reporting is in the form of a relation between the perceiver and the perceived where the latter are extra-mental objects or conditions of one's environment. It is also not surprising that perception typically yields conceptual categorization of the perceived object or condition: the typical result of a perceptual process is bringing the perceived object under a concept. Genuine perception thus puts the premium in the perceived object, not in the perceptual activity itself or in the perceptual mental state whereby one is typically brought into epistemic contact with one's extra-mental environment. If sentences reporting pain in body parts don't follow the pattern of sentences such as (1), i.e., if they are not to be construed as reports of perceptual relations between the perceiver and the perceived, then pain reports are prima facie not perceptual reports, reports to the effect that one stands in a perceptual relation to something extra-mental.

What do sentences such as (2) then report, if not the obtaining of a perceptual relation between the perceiver and a perceived object or condition? The answer is anti-climactic: why, of course, they report pain experiences, i.e., mental states or events with certain perhaps complex but characteristic phenomenal character or content. Even though this may be anti-climactic in the sense that it's almost a truism given our ordinary notion of pain, it is not enough to remove puzzle. Experiences are in the head, and for most physicalists they are in the head by being realized in the brain or central nervous system. If sentences like (2) are reports of the occurrence of certain kinds of experiences, we still need to understand what is going on when we seem to locate pains in body parts. Obviously, if pains are experiences, locating a jabbing pain in the back of my hand is at best confused. How do we reconcile this with the fact that I say something true when I utter (2) (even when, let's assume, there is nothing physically wrong with my hand — say, because I am suffering from a centrally caused chronic condition, which is not uncommon)?<sup>2</sup> In other words, how do we reconcile the common sense understanding of pain as a subjective experience with the comfort and ease with which the very same common sense routinely attributes pain to body parts?

There is no easy and comfortable answer to this question. The common answer on the part of the perceptual and strong representational theorists such as Pitcher, Armstrong and Tye, which I share, is to say that sentences like (2) report the occurrence of experiences that represent a certain condition (or constellation of conditions) in one's

---

concept of pain, they all follow the same semantic pattern in terms of not involving attribution of physical properties or conditions.

<sup>2</sup> Indeed, chronic pain syndromes are not restricted to rare cases like phantom limb pains and referred pains (although the latter are more common than the former). There are more than 1500 pain clinics in the US alone mostly devoted to treat chronic pains, almost all of which are centrally caused pains felt in bodily locations that are not in any pathological conditions. 40% of all Americans suffer from chronic pain at least once and usually late in their lives. Indeed as I write this, I am all too painfully aware of my sciatic pain that I have been suffering for the last two years. I know well that my right leg is just fine as I am told by many medical experts — the cause of the pain is the pressure put by a bone perturbation on the sensory nerve coming out of my lower spine.

body regions. What they represent exactly is a matter of controversy, but Tye's answer is that pain experiences represent tissue damage (or, let's add, an objective physical condition conducive to tissue damage if sustained) — so their location is the location of the actual or potential tissue damage. The experiences represent tissue damage nonconceptually in a way analogous to how our visual experiences represent colors of surfaces. If colors are objective physical conditions of surfaces, the way our visual system represent these conditions is such that we can't necessarily conceptualize them as such. So it is no objection to this view that we don't conceptualize what pain experiences represent as tissue damage. How then do we conceptualize them on the basis of how pain experiences represent tissue damage? The confusion and the error found in the common sense understanding of the truth conditions of sentences like (2) stems from the fact that we seem to use the same term (and perhaps the same concept pre-reflectively) for the damaged condition of bodily tissues as the term we use to report their experiences, namely pain. It might be helpful to distinguish the former use of 'pain' with a subscript, 'pain<sub>td</sub>', reflecting that it refers to the actual or potential tissue damage, to a physical condition of the tissue which is typically the cause of the pain experience, i.e., pain<sub>e</sub>.<sup>3</sup>

Armstrong summarizes the resulting account of the semantics of sentences like (2) thus:

---

<sup>3</sup> Tye draws a similar distinction, but he adds a puzzling qualification: "The term 'pain', in one usage, applies to the experience; in another, it applies to the quality represented insofar as (and only insofar as) it is within the content of a pain experience. Which quality (or type) is represented? Pain experiences normally track tissue damage. So, tissue damage is the obvious naturalistic candidate for the relevant quality." (this volume, p. ??, my emphasis). This qualification is puzzling in the light of the identification of this quality with tissue damage. It seems to come down to this: the term 'pain' in the second ordinary usage applies to a relation, to the tissue damage as being experienced by its "owner." So the tissue damage is not a pain in any sense, but it's what we might call a proto-pain (cf. Dretske (2003, ms.)), a proto-pain which becomes a pain, in the second sense, when and only when actually being experienced by someone (i.e., actually being felt or experientially represented by its owner through the appropriate somatosensory modality). If this is what it comes to, then it seems that it is incorrect to claim that there is an ordinary usage of pain with this sense — if the claim is not simply a terminological stipulation or is meant to be part of a substantive theory of what pains are. For to say that there is such a sense is simply to claim that sentences like (2) has an ordinary reading similar to that of sentences like (1), according to which by uttering (2) we sometimes ordinarily mean something that implies that

- (3) there is a tissue damage in the back of one's hand such that one is feeling it (by somatosensorily experiencing it).

I doubt there is such an ordinary sense of 'pain' — see below.

I am also puzzled about the way Tye describes what a pain experience represents, when he says (in the paragraph preceding the above quotation where he seems to be making the same claim) that pain is "also a certain quality or type insofar as (and only insofar as) that quality or type is experientially represented." Assuming representationalism about experiences is right, experiences don't just represent a quality or type, they represent them as instantiated at a time and place (even if vaguely). Unlike concepts, that is the only way they represent qualities or types — and that is for good reason again: having a sensory experience is having a state that purports to carry information in the technical sense of the term. I can entertain a concept in the abstract without attributing it to anything. But I cannot entertain experiences like that (by the way, sensory imagination may be like entertaining a concept, but it is certainly not experiencing). Having an experience that purports to carry information that a certain quality is instantiated at a place at a time is having an illusion or hallucination if it turns out that there is no such quality actually instantiated there and then.

'I have a pain in my hand' may be rendered somewhat as follows: 'It feels to me that a certain sort of disturbance is occurring in my hand, a perception that evokes in me the peremptory desire that the perception should cease' ...

The force of the word 'feels' in this formula is no more and no less than the force of the word 'feels' in 'My hand feels hot', where this latter sentence is so used that it neither asserts nor excludes my hand being hot in physical reality... Now while there is no distinction between felt pain and physical pain, there is a distinction between feeling that there is a certain sort of disturbance in the hand, and there actually being such a difference. Normally, of course, the place where there feels to be such a disturbance is a place where there actually is a disturbance, but in unusual cases, such as that of the 'phantom limb', or cases of 'referred pain', there feels to be disturbance in a place where there is no such disturbance.

The problem of the 'location' of physical pains is therefore solved in exactly the same way as that of the location of 'transitive' sensations, such sensations of pressure... [T]he 'location' of the pain is ... an intentional location. There feels to be a disturbance of a certain sort in the hand, whether or not there is actually such a disturbance.

This account of the location of pain enables us to resolve a troublesome dilemma. Consider the following two statements: 'The pain is in my hand' and 'The pain is in my mind'. Ordinary usage makes us want to assent to the first, while a moment's philosophical reflection makes us want to assent to the second. yet they seem to be in conflict with each other. But once we see that the location of the pain in the hand is an intentional location, that is, that it is simply the place where a disturbance feels to be, but need not actually be, it is clear that the two statements are perfectly compatible. (Armstrong 1968, pp. 314–6)

So what sentences like (2) in fact attribute in the first instance is not pain<sub>td</sub> but rather pain<sub>e</sub>. What I do when I correctly utter (2) is to self-attribute an experience which then attributes a physical condition to a location in my hand. Intuitively, by uttering (2) I am saying something like "I am undergoing an experience which tells me that something untoward going on in the back of my hand." If so, the fact that nothing untoward going on in my hand (nothing physically wrong with my hand) doesn't make (2) false. The fact that I can still correctly point to where it really hurts in my hand after hearing from my doctor that nothing is wrong with my hand is explained by reinterpreting what I say and do with that gesture: I am still undergoing an experience which represents my hand as something wrong with it.

So according to this perceptualist/representationalist account, when we make claims about where it hurts (attribute pain to bodily locations), we in fact rescind from committing ourselves to there being anything physically wrong in those locations — even though we normally expect to find some physical disorder in them. In other words, when our attention is drawn to those locations and the felt qualities in those locations, we don't strictly speaking commit ourselves to find anything in them. What is primary, then, is the experience itself, i.e., that we are undergoing a pain experience in the first instance. Only secondarily is our attention drawn to what those experiences point to, and that only in a non-committal way. Note that, on the proposed analysis of (2), to say that my pain experience attributes a physical condition to a bodily location is not to say that I do so attribute it. On the contrary, the fact that I leave open whether the attribution is correct shows that I don't. I simply report an experience that tells me something, whether or not I come to believe what it tells me is a matter of factors not to be read into the analysis of what the truth-conditions of (2) are.

Ironically, this representationalist treatment of pain reports is in tension with the general thrust of any perceptual/representational account of experience, because what it reveals is that in pain our immediate interest (epistemic or practical) is in the experiences themselves in the first place, rather than in what objects or conditions these experiences represent. As we noted above, it is the other way around in standard perception — genuine perception where the concepts we are induced to apply apply in the first instance to the objects and conditions that the inducing perceptual experiences represent. We see an apple and identify it as such. We see a discoloration and identify it as such. In pain, what we identify in the first instance as pain is the experience itself — even though our ordinary reports confusingly point to what these experiences represents (assuming the above account). Indeed, the appropriate classification of access to pain makes it a form of introspection, rather than exteroception or perception of extra-mental environment. Even when we attend to the qualities in a body region that our pain experience represents, the appropriate classification of our activity is introspection. Indeed, Tye himself, repeatedly uses ‘introspection’ and talks about introspecting the attended qualities even when he argues for the transparency of pain experiences and identifies these qualities with tissue damage and its features. The term ‘introspection’ is appropriate given what we have said so far, but it is puzzling to see that it’s Tye’s choice of word — because whatever else introspection may be, I take it that it is a conceptual truth that introspection is a method of access to mental phenomena and not a method of access, like standard perception, to extra-mental environment. Whatever else the transparency of genuine perceptual experiences involves it must take us to the extra-mental world in a committal way — or roughly equivalently: sentences like (1) must admit an analysis according to which perception is a relation (at least) between the perceiver and the perceived extra-mental objects or conditions.<sup>4</sup> But what we access in pain is the experience itself and what it represents without a commitment to its veridicality. If pain experiences are a species of genuine perceptual experience whose phenomenal content is wholly constituted by its representational content as strong representationalism demands, then it’s a mystery why pain experiences and the conceptual categorization they induce work the way they do, quite differently from standard perception.

---

<sup>4</sup> It is a little puzzling why Tye thinks that pain experiences are transparent. He correctly points out where our attention is drawn in feeling pain in a bodily location and emphasizes that our behavior also favors the location we feel pain at. But why would anyone be inclined to think that these show the kind of strong transparency of pain experiences Tye favors? Sense-datum theorists, for instance, would have no difficulty in explaining these facts even though they would reject the strong transparency of pain experiences — see Jackson (1977) among others. When we feel pain in a body location L, we quite naturally expect to find something physically wrong there, call the general condition TD that we expect to find in L — whatever it may turn out to be. If our feeling pain in L were transparent, if it were simply a matter of perceiving TD in L, we would not expect to find TD there, we would find TD there. In cases where we don’t find TD in L, we don’t rescind from reporting pain in L. In such cases, we still correctly report pain in L. Compare visual hallucinations. I report a pink butterfly flying ahead of me at three o’clock position. My report is based on a visual hallucination. My report is false. Not only that, my report that I see a pink butterfly flying ahead at three o’clock is also false. Upon realization and only then do I revert to introspective mode and report correctly that I visually seem or seemed to see a butterfly. Nothing of this kind happens in reporting pain. If we find that there is no TD in L, we do not make any correction, we still keep reporting the situation, correctly, as exactly in the same way as we did the first time. This reveals that a pain report is an introspective report from the very beginning, hence not a perceptual report in the first place.

This was the worry I raised in my (2001) to which Tye reacts in his essay in this volume by saying that the worry was misplaced. In the following two full paragraphs Tye attempts to explain why, I couldn't discern an explanation of why the worry was indeed misplaced. Tye writes:

The worry is misplaced. When we introspect our color experiences, we issue reports about the apparent colors of things. Here the information flow goes from the awareness of the color red to the application of the concept EXPERIENCE OF RED or the concept LOOKS RED. By being aware of red, I am aware that I am having an experience of red. Correspondingly, when we introspect our pain experiences, we are aware of (among other things) tissue damage and by being aware of those things, we are aware that we are in pain. For these cases, the information flow is of just the same type. Normally, awareness of red is caused by red. Normally awareness of tissue damage is caused by tissue damage. Both awarenesses of something external (to the head) bring about an awareness that an experience of a certain sort is present. (2005, this volume)

From this, it's easy to get the impression that what happens in perceiving red is exactly the same as perceiving tissue damage by feeling it, and the introspective processes involved in both are thus exactly the same. Note that when I perceive red and identify it as such by the application of RED, I don't do so by first introspecting my visual experience representing red. On the contrary, on Tye's view, introspection of visual experience of red requires that one have the concept of red (and the concept of experience); in other words, one must already have the capacity to apply RED to red things via one's experience of red. As Tye says, "introspection of experiences ... is a reliable process that takes awareness of qualities represented by the experiences as input and yields awareness that a certain kind of experience is present as output" (ibid.). What kind? The kind that represents those qualities. Since the output is a thought, it requires the concepts that represent what the experience represents. Hence you need the concept RED to introspect your visual experiences of red. Since, according to strong representationalism, the phenomenal content of an experience is completely exhausted by its representational content, introspecting the phenomenal content of an experience is tantamount to have a thought as the output of that reliable process about what the experience represents qua experience. If you don't have the concepts of the proprietary range of qualities represented by an experience, you are not only conceptually blind to those qualities, you are also blind to the experiences representing those qualities. This is true in all introspection according to Tye's representationalism. Experiencing red is a matter of perceiving red and conceptually identifying it as such in the first place. Its introspection comes later — if at all.

But with pain, it's the other way around. Tye says "when we introspect our pain experiences, we are aware of (among other things) tissue damage and by being aware of those things, we are aware that we are in pain." I don't think this is correct. When we introspect our pain experiences, we are not aware of tissue damage in the first place. For one thing, tissue damage is not the kind of thing that we can introspect. For another, when we come to notice pain and report it, we come to notice and report the occurrence of an experience in the first place — as Tye himself says: "when I am introspectively aware that I am in pain, the concept I apply is the concept PAIN and not the concept TISSUE DAMAGE. I agree also that the concept PAIN applies to the experience I undergo" (p. ??, this volume).

This being so, when we introspect our pain experiences, we are aware of our experiences in the first place. The first concepts that we are immediately prompted by the pain experience to apply are concepts that apply to that experience (i.e., to pain), not to what the experience represents, i.e., to tissue damage and its features — as indicated by the lack of our commitment to the veridicality of what pain experiences represent (assuming that they represent tissue damage) when we report pain in body parts. It is all too easy to be misled by our confused common practice of attributing pain to body parts that such practices, all by themselves, support strong representationalism or somehow show that pain experiences are transparent in the strong metaphysical sense that Tye intends it. But they show no such thing — without additional assumptions or arguments.

So if “pain perception” is a matter of introspection in the first place despite being sustained by the same kind of information flow mechanism underlying other genuine perceptual processes, it is legitimate to express worry and ask why pain processing doesn’t conform to the norms of standard perceptual processing. In other words, why is pain a matter of introspection of an experience in the first place, and only then a perception of tissue damage — if at all? I am yet to see why this worry is misplaced.

All the early (direct realist) perceptual theorists such as Pitcher (1970, 1971) and Armstrong (1962, 1968) among others recognized the difficulty as a puzzle for perceptual theories of pain, and acknowledged that this puzzling asymmetry between pain and standard perception needs an adequate explanation consistent with their direct realist perceptualism.<sup>5</sup> And they offered an explanation. Their basic idea was that although pain experiences are indeed genuine perceptual experiences, because of their immediate negative affect, the informational value they have about tissue damage is shadowed by a more “peremptory desire that the perception should cease,” a desire evoked by those experiences, and thus forcing the cognitive/conceptual system to adjust its utilities by anchoring the locus of immediate conceptual identification to the pain experiences themselves, rather than to what these experiences represent or carry information about, i.e., tissue damage — as reflected in our ordinary concept of pain and pain reports. Whether or not this is an adequate explanation is something I cannot discuss in this essay.<sup>6</sup> The point is that there is a legitimate worry here and an adequate representationalist explanation is needed.

It is important to note that I am not objecting to the claim that pain experiences are (or, involve) perceptual experiences or that their phenomenal content is constituted by their representational content. None of what I have said above shows that these claims are false. In fact I think that pain experiences are genuinely perceptual (as well as affective/emotional) and their phenomenology (except their affective phenomenology — see below) is exhausted by their representational content. Indeed I

---

<sup>5</sup> Armstrong writes for instance: “Now it may well be asked why, in this important field of perception [intransitive bodily sensations], ordinary discourse restricts itself to sense-impression statements, instead of making perceptual claims about what is actually going on in our body. Several reasons may be advanced. In the first place, we have seen that the [conative] reactions characteristically evoked by these bodily impressions are determined almost solely by the impression, whether or not it corresponds to physical reality. Since the reactions are an important part of our concept of almost all the intransitive bodily sensations, it is natural to talk about what feels to us to be the case, whether or not it is the case” (Armstrong 1962, p. 125). Cf. Pitcher (1970, p. 385) — among other places.

<sup>6</sup> In Aydede (forthcoming) and Aydede & Güzeldere (forthcoming) I argue that this sort of explanation is not adequate — at a minimum incomplete. In the second essay, we offer a fuller account in the context of an information-theoretic account of introspection.

favor the same kind of representationalist account of pain attribution sentences as Tye does and other perceptual theorists before him did. But what is needed is an adequate account of why there is a such marked difference between obviously perceptual processes (such as vision, hearing, etc.) and pain (as well as other intransitive bodily sensations like itches, tickles, tingles, orgasms, etc.) — especially when (direct realist) perceptualism and strong representationalism about pain crucially predicts otherwise.

Finally, I want to comment briefly on Tye's claim that his strong representationalism can capture the negative affective phenomenology of normal pain experiences. In my (2001), I criticized Tye's brief treatment of pain affect that he had presented in his (1996, 1997). My criticism was that being a pure qualia representationalist Tye could not account for the negative affect of pain experiences in purely representationalist terms and that he was forced to adopt what I called cognitivism about negative affect, which was also the option adopted by early perceptual theorists such as Armstrong and Pitcher. According to cognitivism, the affective qualities of pain experiences are somehow constituted by one's cognitive conative reactions to them, by pain's evoking in one "the peremptory desire that the perception should cease." (Armstrong 1968, p. 314) In his new essay in this volume, Tye distances himself from any such way out and proposes that the affective qualitative content of pain is entirely representational as his pure qualia representationalism demands:

People in pain try to get rid of it or to diminish it. Why? The answer surely is because pain feels unpleasant or bad, because it is experienced as such. But what exactly is experienced as unpleasant? One's attention, when one feels pain, goes to a place different from the one in which the experience of pain is located. The qualities that are experienced as unpleasant are located in the bodily location to which one attends (in normal circumstances). People whose pains lack the affective dimension undergo purely sensory, non-evaluative representations of tissue damage of one sort or another in a localized bodily region. Those whose pains are normal experience the same qualities, but now those qualities are experienced by them as unpleasant. It is precisely because the qualities are experienced as unpleasant or bad that people have the cognitive reactions to them they do, reactions such as desiring to stop the pain.

To experience tissue damage as bad is to undergo an experience which represents that damage as bad. Accordingly, in my view, the affective dimension of pain is as much a part of the representational content of pain as the sensory dimension is. (2005, this volume)

So the experience I undergo when I feel a jabbing pain in the back of my hand feels the way it does because it represents tissue damage<sup>7</sup> in my hand and it represents it as bad. In other words, it represents tissue damage as having the quality of being bad.

Of course, normally having one's tissue damaged is bad. We also normally think/judge that it is bad. Not only that, experiencing tissue damage (i.e., feeling pain) is bad. We think/judge so too. But none of these truisms (all by themselves) can support the claim that the experience represents the damage as bad. And so they should not be confused with this last claim. Indeed what does it mean for one's experience to represent tissue damage as bad? What is the quality of the tissue damage itself that is detected or tracked by the experience so that we can say the experience

---

<sup>7</sup> In a PANIC way — see his (1996). Presumably my experience represents a myriad of other features — let's simplify and just focus on tissue damage and its badness in what follows.

carries information about it? What objective quality of tissue damage could be such that its possession by the tissue damage can be detected by sensory experiences such as pain and turn them also into affective/emotional experiences? Note that Tye is eager to insist that sensory experiences has a proper range of sensory qualities they track. It is the representation of these qualities that gives the experiences their peculiar phenomenal content. (This comes out especially clearly in Tye's response to Maund's objection that things can look rusty, wooden, metallic, etc.) Now we learn that among these proprietary qualities our experiences can detect or track are the badness or goodness of those very same qualities. Recall that Tye presents himself as an informational theorist when it comes to the metaphysics of psychosemantics (indeed this is what the "I" in "PANIC" captures).

I am not sure how to understand Tye's proposal. It seems that whatever the badness or goodness of those qualities properly detected by experiences might come to at the end, these second order qualities are simply not the kind of qualities that can be detected or tracked (in the technical information-theoretic sense). One reason for thinking this is Tye's own test: there could be a molecule-by-molecule identical tissue damage which is not bad (not just experienced as bad, but just not bad — if we abstract away from the regular connotations of 'damage' and think of it as composed of whatever physical features and their configurations constitute the damage). There doesn't seem to be any natural property of a tissue damage simple and suitable enough to allow itself to be transduced. Surely the burden of proof here is on Tye to show us that pain experiences normally carry information about a certain quality of tissue damage that constitutes its badness. What natural property is that so it can be detected? If it's not a natural property, what reasons are there to think that it's the kind of property whose instantiations can be detected in the information-theoretic sense? We need to be told.

Perhaps, Tye would want to relax his metaphysics of psychosemantics and allow that an experience can come to represent something as bad not in virtue of some sort of informational relation, but rather in virtue of its causal/functional role.<sup>8</sup> Because these experiences play the role they do in the complex mental and behavioral economy of creatures capable of feeling pain, they come to represent the tissue damage they detect as bad. So the sensory phenomenology of pains gets explained by an informational psychosemantics and their affective-emotional phenomenology by a functional role psychosemantics. I think there is something to this proposal. But note that if Tye opts for this way out, then representationalism about affect becomes otiose. It is a needlessly complex story: the functional role of pain experiences makes for a certain (valuational) representational content which is then to be reductively identified with the affective phenomenology of pain experiences. Why not simply say that a certain psychofunctional role that pain experiences play itself constitutes their affective phenomenology?<sup>9</sup> Indeed, functionalism or psychofunctionalism in the philosophy of mind has always treated pains as paradigm examples of mental states whose phenomenology can be plausibly captured by functionalist proposals. What made this *prima facie* plausible was pain's affective phenomenology (as opposed to its sensory

---

<sup>8</sup> I haven't seen Tye taking this route before. So my guess in the text is purely speculative. But I hope he would find it plausible and consider it for adoption..

<sup>9</sup> I defended this kind of representational-cum-psychofunctional account of affectively non-neutral experiences in my (2000), which also found its way to my (2001). For a more elaborate criticism of pure qualia representationalism along these lines, see my (ms.).

phenomenology), which is essentially connected with its being an intrinsically motivational/drive state. But once psychofunctionalism is allowed, we don't need to make a detour through representationalism. In fact, it is hard to see how an experience's representing something as bad can be an intrinsic motivator all by itself. To explain how, a representationalist has to advert to additional (learning?) mechanisms to connect representational content with drive and motivation.

On Tye's view even though we don't need to have any concept to have a pain experience with all its normally awful affective phenomenology, we need to have concepts to introspect it and its affective qualities. Recall that according to his official view, one can introspect phenomenal qualities only if one has the concepts denoting what those qualities represent. So the official story about how we come to introspect our pain experiences and their affective qualities must go something like this. Simplifying quite a bit: we are capable of having pain experiences representing tissue damage and representing it as bad; we are capable of applying concepts like PAIN<sub>td</sub> and BAD (or their equivalent) to the tissue damage and to its badness. We must also be capable of deploying concepts like EXPERIENCE and REPRESENTS (or their equivalent). When all these capacities are in place, our reliable introspection module can be activated to output an introspective judgment to the effect that we experience pain<sub>td</sub> as bad and know what it is like to have experiences with that content. I am in sympathy with this kind of view of how introspection works when it comes to how we introspect affectively neutral standard exteroceptive experiences.<sup>10</sup> But it's simply not credible that this is the way we come to introspect pain and its qualities (as well as all the other intransitive bodily sensations). Pain, as already discussed, is a form of introspection from the very beginning, not a form of perception in the first place.

Pain and other intransitive bodily sensations pose serious difficulties for (direct realist) perceptual theorists and for strong representationalists. What those of us with naturalistic and representationalist convictions need to understand is that bodily sensations are special: instead of trying to remove every bit of difficulty in the way of smoothly and completely assimilating them to other obviously perceptual processes, we need to make room for a broader but messier conception of perception which would involve bodily sensations with all their peculiarities. In other words, we need to pay close attention to how bodily sensations work if we are ever going to succeed in understanding how, despite all their peculiarities, they could after all turn out to be (or at least, involve) genuine perception.

---

<sup>10</sup> Indeed, in Aydede & Güzeldere (forthcoming), we proposed a detailed empirical (but speculative) theory about how we introspect our exteroceptive experiences that relies on information-theoretic principles laid out in Dretske's seminal (1981) work, which, we believe, is in broad agreement with Tye's views.

## REFERENCES

- Armstrong, D. M. (1962). *Bodily Sensations*. London: Routledge and Kegan Paul.
- Armstrong, D. M. (1968). *A Materialist Theory of the Mind*. New York: Humanities Press.
- Aydede, M. (2000). "An Analysis of Pleasure vis-à-vis Pain." *Philosophy and Phenomenological Research*, 61(3): 537–570.
- Aydede, M. (2001). "Naturalism, Introspection and Direct Realism about Pain." *Consciousness & Emotion*, 2(1): 29–73.
- Aydede, M. (forthcoming). Pain. *The Stanford Encyclopedia of Philosophy (Fall 2004 Edition)*, E. N. Zalta (Ed.):  
URL=<<http://web.clas.ufl.edu/users/maydede/pain/pain.sep.pdf>>
- Aydede, M. (ms.). "Against Pure Representationalism about Qualia: The Case of Intransitive Bodily Sensations." University of Florida, Department of Philosophy, 2003: URL=<<http://web.clas.ufl.edu/users/maydede/AgainstPureRep.pdf>>
- Aydede, M., and G. Güzeldere (forthcoming). "Cognitive Architecture, Concepts, and Introspection: An Information-Theoretic Solution to the Problem of Phenomenal Consciousness." *Noûs*:  
URL=<<http://web.clas.ufl.edu/users/maydede/introspection.pdf>>
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, Massachusetts: MIT Press.
- Dretske, F. (2003). How Do You Know You Are Not A Zombie? *Privileged Access: Philosophical Accounts of Self-Knowledge*, B. Gertler (Ed.), Hampshire, UK: Ashgate Publishing.
- Dretske, F. (ms.). Knowing It Hurts.
- Jackson, F. (1977). *Perception: A Representative Theory*. Cambridge: Cambridge University Press.
- Pitcher, G. (1970). "Pain Perception." *The Philosophical Review*, 79(3): 368–393.
- Pitcher, G. (1971). *A Theory of Perception*. Princeton, NJ: Princeton University Press.
- Tye, M. (1996). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, Massachusetts: MIT Press.
- Tye, M. (1997). A Representational Theory of Pains and their Phenomenal Character. *The Nature of Consciousness: Philosophical Debates*, N. Block, O. Flanagan and G. Güzeldere (Eds.), Cambridge, Massachusetts: MIT Press.
- Tye, M. (2005, this volume). "Another Look at Representationalism and Pain." *Consciousness & Emotion*, 6.