ANOTHER LOOK AT REPRESENTATIONALISM ABOUT PAIN

There is something it is like to feel a sharp pain in the elbow, a migraine headache, a throbbing pain in the leg, a squeezing pain in the heart. These experiences are phenomenologically different, but they also have a common phenomenological core. Do they also have representational content? Some philosophers have thought not. For example, Colin McGinn remarks:

> . . . bodily sensations do not have an intentional object in the way that perceptual experiences do. We distinguish between a visual experience and what it is an experience of; but we do not make this distinction in respect of pains. Or again, visual experiences represent the world as being a certain way, but pains have no such representational content. (McGinn 1982, p. 8)

McGinn's comments on visual experience are uncontroversial. Visual experiences have correctness conditions. For the subject of the experience, the world *seems* a certain way, the way represented by the experience. The way the world is represented is the representational content of the experience. The experience is accurate if the world is that way, inaccurate otherwise.

McGinn's denial that pain experiences have representational content is open to dispute, however. Some philosophers allow that pain experiences have such content as far as the bodily location of pain goes, but they deny that this content, which is shared with other bodily sensations, has anything to do with the characteristic phenomenology of pain experiences (Block 1996). Still other philosophers maintain that pain experiences have a distinctive representational content—for example, they all represent tissue damage at a given bodily location—but they deny that this content captures their phenomenology (Chalmers 1996).

My own view is that pain experiences have a distinctive representational content and that this content *is* their phenomenal character. This view, which is sometimes known as "strong representationalism" for pain, has come under attack from several different quarters in the recent literature. The purpose of this paper is to lay out the strong representationalist position for the case of pain in more detail than has been done before (see Tye 1995, 1995b, 2000 for earlier discussions) and to take up various criticisms along the way.

I: Some preliminary remarks on pains as representations

You cannot feel my pains and I cannot feel yours. Even if we are Siamese twins, joined at the hip and stung there by a bee, intuitively there are two pains, yours and mine. Maybe your pain is a bit different from mine. Without further information, that's an open question,, which it wouldn't be, of course, if there were just one pain.

Pains cannot exist without owners. There cannot be a pain with no-one or no creature around to feel the pain. To be sure, there is the phenomenon of unnoticed pain, but even an unnoticed pain has an owner — the person who fails to notice it.

That pains are necessarily private and necessarily owned is part of our folk conception of pain and it requires explanation. The obvious explanation is that pain is a feeling or an experience of a certain sort. That is certainly how scientists think of pain. Witness the definition of pain by the *International Association for the Study of Pain*, as "an unpleasant sensory and emotional experience associated with actual or potential tissue damage, or described in terms of such damage" (in the journal, *Pain*, 1986). And it is also part of our commonsense conception.

Since you cannot feel my feelings any more than you can laugh my laughs or scream my screams, and since there cannot be a feeling without a subject for the feeling any more than there can be an unlaughed laugh or an unscreamed scream, pains, conceived of as feelings, are necessarily private and necessarily owned.

Pains, viewed as experiences, intuitively can be misleading or inaccurate. Take the case of phantom limb pains. You have no right leg and yet you experience pain in the leg. How is that possible? Answer: in undergoing the pain experience, you *hallucinate* a right leg.

Referred pains also are nonveridical. You can feel a pain in the left arm, when there is nothing wrong with the arm, the cause of the experience being a disturbance in the heart. Such a pain intuitively is inaccurate or misleading; for without additional information, on the basis of the pain, you would be disposed to nurse the arm, to rub it, to believe that something is awry in the arm itself. The case is one of *illusion*.

If pains are representations, what do they represent? The obvious answer is pain. After all, if I feel a pain in a leg, and I attend to *what* I am feeling, I attend to a quality that *seems* to be tokened in my leg, a quality that I strongly dislike and that, in one ordinary sense of the term 'pain,' is surely pain. A theory that denies this categorically is at odds with what we ordinarily believe.

But how can this claim be accommodated? The answer is to acknowledge that pain is not only a certain sort of experience but also a certain quality or type *insofar as* (and only insofar as) that quality or type is experientially represented.

So, the term 'pain' does double duty (Harman 1990). A pain in a leg (viewed as an experience) represents that a certain quality is tokened in the leg. It is accurate if that quality is tokened there; inaccurate otherwise. 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 is not to say that pains always are alike with respect to the kind of tissue damage represented. As I noted in my 1995, a twinge of pain represents a mild, brief case of damage. A throbbing pain represents a rapidly pulsing disorder. Aches represent regions of damage inside the body rather than on the surface. These regions are represented as having volume, as gradually beginning and ending, as increasing in severity, as slowly fading

away.[1] The volumes so represented are not represented as precise or sharply bounded. This is why aches are not felt to have precise locations, unlike pricking pains, for example. A stabbing pain is one that represents sudden damage over a particular well-defined bodily region. This region is represented as having volume (rather than being two-dimensional), as being the shape of something sharp-edged and pointed (like that of a dagger).[2] In the case of a pricking pain, the relevant damage is represented as having a sudden beginning and ending on the surface or just below, and as covering a very tiny area. A racking pain is one that represents that the damage involves the stretching of internal body parts (e.g., muscles).

II. On exportation and substitutivity in pain contexts

It is well known that in belief contexts, co-referential terms cannot safely be substituted. For example,

The ancients believed that Hesperus = Hesperus

is true, while

The ancients believed that Hesperus = Phosphorus

is false. Exportation also fails. For example,

Many children believe that Santa Claus lives at the North Pole

is true, but

Santa Claus is such that many children believe that he lives at the North Pole

is false. Some philosophers have held that these two marks of representation always go together (Quine 1960). But there is no a priori reason why they *should* always go hand in hand; and for pain statements, it is clear that there are cases in which the principle of substitutivity holds but exportation fails. For example,

(1) Desmond feels a pain in his finger

together with

(2) Desmond's finger is the body part he recently cut

entails

(3) Desmond feels a pain in the body part he recently cut.

Here the substitution of 'the body part he recently cut' for 'his finger' is safe, as are all substitutions in analogous inferences. But, given the possibility of phantom limb pain, (1) does not entail

(4) Desmond's finger is such that he feels a pain in it.

---

[1] Cp. Armstrong 1962.

[2] I am not suggesting here that one cannot have a stabbing pain unless one has the concept of a dagger.

Exportation intuitively also fails with respect to the expression 'a pain' in (1). To see this, suppose that (1) is true. Still

(5) A pain is such that Desmond feels it in his finger

may well be false. For Desmond does not feel a *feeling of pain* in his finger; he feels *a pain* in his finger; and if the case is one of phantom limb pain, it seems clear that the only genuine token pain that exists is Desmond's feeling (or experience) of pain.

Consider now the following statement similar to (1) but in which 'pain' is not part of an indefinite description:

(6) Samantha feels pain in her toe.

Can we safely substitute 'tissue damage' for 'pain' here so that if (6) is true then

(7) Samantha feels tissue damage in her toe

is true also? Well, if the principle of substitutivity holds, we can so substitute if

(8) Pain = tissue damage

is true. But (8) is not true, on the view I am proposing. For my claim is that 'pain,' in one sense of the term 'pain,' applies to tissue damage (on the assumption that pain experiences represent tissue damage) only insofar as it is within the content of a pain experience. This highly qualified claim does not license the unqualified identity claim made by (8).

The fact that exportation fails in statements such as (1) shows pain experiences are indeed representational; for how else can the failure of exportation be accounted for? The fact that there is no failure of substitutivity in (1), even though exportation fails, suggests that pain experiences have representational contents different in kind from beliefs. Belief (and thought) contents individuate in a more fine-grained way than pain contents. On the natural assumption that conceptual contents generally individuate in the same manner as belief contents, it follows that pain experiences have nonconceptual representational contents. Furthermore, these contents are nonconceptual not just in the sense that the subjects of pain experiences need not have the concepts used to state their correctness conditions; for contents that are nonconceptual in this sense *can* be the contents of beliefs (though not beliefs the subjects of the experiences need be in a position to have themselves). Pain experiences have nonconceptual contents that are *different in kind* from belief contents. I shall return to the topic of nonconceptual content later.

III. The location of a pain

We locate pains in bodily parts. We say things like "I have a dull pain in my right leg" and we use descriptions for pain that appeal to bodily parts, for example, "the burning pain in my stomache." The *feeling* of pain isn't in my right leg or my stomache, but the pain, it seems, is. How can this be?

There are four possible strategies for handling pain location. One is to adopt a representational account. More on this shortly. A second is to maintain that pains are real objects of the experience of pain. When I feel a pain in my chest, for example, the pain really is in my chest even if the feeling of pain is not. That, however, cannot be right. The pain, whatever else it is on such a proposal, is a mental object for the feeling. And mental objects cannot exist in chests or legs any more than such objects can exist in walls or tables. Of course, I can feel a pain in a hand (that has recently been amputated) and I can place the stump of my arm against a wall or a table. But the pain I feel isn't *in* the wall or table!

A third response is to say that the word 'in' in such contexts has a special causal sense. Pain and the feeling of pain are one. In conceiving of a pain as being in a leg, we are conceiving of it as being caused by damage to the leg, not as being spatially inside the leg. That again seems very implausible. For one thing, it is ad hoc. For another, unless it can be shown to be the only alternative that works, it unnecessarily multiplies word senses. Why insist that 'in' here means something different from what it means when it is used elsewhere?

Those who take this view try to motivate the proposal by noting that inferences like the following are invalid:

(9) The pain is in my fingertip

(10) The fingertip is in my mouth

therefore,

(11) The pain is in my mouth.

How could this be, they ask, if the word 'in' is used in (9) and (10) in its ordinary spatial sense (Jackson 1977; Block, 1983)? After all, if *A* is spatially inside *B* and *B* is spatially inside *C*, *A* must be spatially inside *C*.

This line of argument would provide some support for the causal proposal, if there were no other explanation of the inference failure consistent with the claim that 'in' in such contexts has an ordinary spatial sense.

One alternative explanation is offered by Paul Noordhof (2001). According to Noordhof, the inference from (9) and (10) to (11) displays "precisely the same invalidity" (p. 96) as the inference

(12) There is a hole in my shoe

(13) The shoe is in the box

therefore,

(14) There is a hole in the box.

What such inferences reveal, Noordhof claims, is only that "there is a sense of 'in' that is used to describe the state of an object" (p. 97).

Unfortunately, Noordhof's claim is opaque. One way to read him is as supposing that talk of holes in things is best understood in terms of talk of the things' being in the state

of being holed (or perforated). However, quantification over holes is unavoidable in ordinary discourse putatively about holes.[3]

Another way to read Noordhof is as claiming that talk of holes in things is to be understood in terms of talk of the things being in the state of having a hole, where the latter state requires for its tokening that there exist a hole. But then what is the proposed reading for (9) or (11)? That my fingertip (or mouth) is in the state of having a pain? That cannot be right.

Furthermore, although the specimen statements of pain location, (9) and (11), arguably require for their truth that the relevant body parts exist, there are many such statements that lack this requirement. For example,

(15) I feel a pain in my fingertip

can be true, even if both my hands have been amputated. In this case, there is no bodily part to be in the relevant state. So, the invalidity of the inference from (15) and (10) to

(16) I feel a pain in my mouth

is left without any satisfactory explanation.

Once the representational approach to pain location is adopted, there is a straightforward account of this inference failure. The inference from (15) and (10) to (16) fails just as does the inference

(17) I seem to see a triangle within a circle

(18) The circle is within a larger square

therefore,

(19) I seem to see a triangle within a larger square.

But what of the inference failure from (9) and (10) to (11)? Noordhof's claim is that it cannot be explained by appeal to representation; for, on the assumption that (9) and (11) require that my fingertip and mouth exist, the relevant pain representations are now *de re* with respect to these items and thus the inference should go through.

Noordhof is mistaken. There is more than one ordinary spatial sense of 'in.' Where there is a hollow physical object, *O*, the claim that something *X* is in *O* can be understood either to assert that *X* is within the cavity bounded by *O* or to assert that *X* is (at least partially) embedded within a portion of the cavity-surround (the top, bottom, and sides of *O*). For example,

(20) Tanya has a ring in her nose

uses the second sense of 'in,' whereas

(21) My car is in the garage

uses the first.

---

3 For compelling arguments here, see Casati and Varzi, 1994. The reason the inference from (13) and (14) to (15) fails is that the hole in my shoe is an entity occupying a fillable discontinuity in the shoe's surface. It is not thereby an entity occupying a fillable discontinuity in the box's surface, even though the shoe is inside the box.

The representationalist explanation for the inference failure from (9) and (10) to (11) is now as follows: the only pain in reality is my experience of pain. According to (9), that pain is accurate if and only if my fingertip is such that there is tissue damage spatially within it (which damage may itself be classified as pain in a second sense of that term insofar as, and only insofar as, it is in the content of my pain experience). (10) asserts that my fingertip is spatially within the cavity bounded by my mouth. From these two claims, it does not follow that my pain experience is accurate if and only if my mouth is such that my pain experience represents that there is tissue damage in a part of its roof, walls, or base. But this is what (11) requires, given that pains represent tissue damage. So, (11), on this proposal, does not follow from (9) and (10).

Perhaps it will be replied that once different spatial senses of 'in' are admitted, that is all that is needed to explain the inference failure. The appeal to representational content is otiose. This misses the point. In the case that (9) and (10) are both true, there is in reality no such thing as a pain in my mouth in *either* spatial sense. Moreover, the experience of pain is in my brain, if it is anywhere, not *in* my mouth. To offer an explanation which supposes otherwise is to offer an explanation based on false assumptions. And that is to offer no explanation at all.[4]

---

4 In a later essay written in response to my 2002, Nordhof (2002) presents further inferences that are intended to show that I must proliferate spatial senses of 'in' over and above the two I distinguished above. This, he suggests, is an unwelcome result for my view. However, I deny that any such proliferation is necessary.

Let us agree that, where *O* is a hollow object, to say that *X* is in *O* is to say either that *X* is within the cavity bounded by *O* (call this 'SENSE 1' of 'in') or that *X* is embedded (at least partly) in the matter making up *O*, i.e., the cavity-surround (call this 'SENSE 2'). Two of the four inferences Noordhof adduces as creating difficulties for me, namely

(B1) There is a flaw in the diamond

(B2) The diamond is in the box

Therefore,

(B3) There is a flaw in the box

and

(D1) There is a fault in my computer

(D2) The computer is in my office

Therefore,

(D3) There is a fault in my office

evidently present no new problem (and it is mystifying as to why Nooordhof thinks that they do). In both cases, a hollow object is involved (a box in one and an office in the other) and 'in' in the conclusion has SENSE 2 whereas 'in' in the second premise has SENSE 1 (and 'in' in the first premise just means 'spatially within'). So, these arguments commit the fallacy of equivocation.

The same is true for the third inference Noordhof cites:

(C1) There is a tremor in my hand

(C2) My hand is in my pocket

Therefore,

IV. Complexities in the content of pain

The view that pain has distinct sensory and affective/emotional components, subserved by different neural mechanisms, was first proposed by Melzack and Casey in 1968, and in the thirty or so years since then it has been shown to be well motivated by both a wealth of clinical data and neuroscientific evidence.[5]

Normally, in a pain experience, both components are present. But in some cases, the affective component is missing. For example, people who undergo prefrontal leukotomies (operations that sever the neural connections in the deep white matter in the frontal lobes) as a last resort for their intractable, constant, severe pain are typically cheerful and relaxed afterwards. They report still having pains, but they no longer *mind* them.

(C3) There is a tremor in my pocket.

The only difference is that, unlike a fault or a flaw, a tremor in the hand has a somewhat vague location. It may occur in a single finger or several without there being anything more definite to say about where the tremor is. It may also occur in the hand as a whole. In this case it is located where the hand is. But the tremor is still spatially inside the hand, allowing that a thing is within a location even if it takes up the entire location (just as a thing is a part of itself). Since a pocket is a hollow object, the inference failure is explained in the same way as the previous two.

Noordhof's own account of these cases is to say that in all of them "an object is described as being in a certain state" (p. 154). But Noordhof denies that he has any reductive ambitions. How then does the appeal to states help? Take the second inference. If my computer is in the state of having a fault in it, and my computer is in my office, it surely *does* follow that my office is in the state of having a fault in it, if 'in', as it occurs within 'state of having an *F* in' means spatially inside. The only solution is either to say, as I do, that 'in' has more than one spatial sense or to introduce further non-spatial senses of 'in'.

Noordhof's remaining inference is

(A1) There is a poisonous gas in the spacecraft

(A2) The spacecraft is in the earth's atmosphere

Therefore,

(A3) There is a poisonous gas in the earth's atmosphere.

This case is a little different from the others. The earth's atmosphere surrounds the earth and other objects not too far away from the earth. In this way, the atmosphere is like a multiple cavity-surround. Something in the atmosphere is either spatially inside it, that is, spatially within what the atmosphere surrounds, or spatially embedded in the atmosphere *qua* surround. In (A1), 'in' has SENSE 1. In (A3), however, assuming we agree that it does not follow from the premises, 'in' has SENSE 2. Given that there is a poisonous gas within the spacecraft and that the spacecraft is within the atmosphere, it does not follow that there is a poisonous gas embedded within the atmosphere *qua* surround.

So, no trouble here for my position and no reason to think that Noordhof has a better explanation in the offing of the invalidity of the pain location inference to which I originally appealed (namely, "I have a pain in my fingertip"; "My fingertip is in my mouth"; therefore, "I have a pain in my mouth.").

5 For useful summaries here, see Melzack and Wall 1983, Price, 1999, 2000.

Similar reports come from people suffering pain who are under hypnotic suggestion or nitrous oxide. Such cases of "reactive disassociation," as Dennett (1978) calls them, are ones in which the distinctive sensory dimension of pain is present but the aversive component is gone. Is pain itself still present? It seems so. The patients *say* that they continue to feel pain. I see no reason not to take these reports at face value. What they show, I suggest, is that pain is not *essentially* an aversive experience.

Pain is essentially a sensory experience, however. Whatever else pain is, at its core, it is a bodily sensation. Take away the characteristic sensory component, and no pain remains. Certain abnormal unpleasant experiences, for example, dysaethesia, lack any sensory component, and they are not classified by their subjects as pains. Other unpleasant experiences, for example, an irritating itch, are sensory, but they aren't pains since the *distinctive* sensory content of pain is missing.

In any event, a *typical* pain experience has both a sensory and an affective dimension. In the earlier sections, I focused on the sensory side of pain. Let me now turn to the affective dimension of pain.

Pain is normally very unpleasant. 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 disturbance that is experienced as unpleasant is 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 sort of disturbance, but now it is experienced by them as unpleasant or bad. It is precisely because this is the case that normal subjects have the cognitive reactions to pain 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.

The representational content of pain, as noted earlier, is nonconceptual. Admittedly, my talk above of the unpleasantness of pain, of its experienced badness may sound cognitive. But I do not intend it to be understood in this way. It seems to me that the most plausible view is that we are hard-wired to experience pain as bad for us from an extremely early age.

Consider the other side of the coin for a moment. A child as young as two months, upon tasting a little chocolate, typically behaves in a way that signifies that he/she wants more. The child will open and close its lips, push forward towards the chocolate, look happy. Why? The answer is that the chocolate tastes good. *That's* why the child wants more. The child's gustatory experience represents a certain taste and the child experiences that taste as good. The taste is experienced as good by the child in that the child undergoes an overall experience which represents the presence of the taste in the mouth and represents it as good.

Intuitively, this is not a cognitive experience. It does not require concepts. For another example, consider orgasm. Orgasm is a bodily sensation, but it is not only that. The most natural description of an orgasm, and indeed of any pleasant experience is "It feels good." One's orgasm represents a certain change in the region of the genitals as good for one, as something apt to benefit, not to harm one.[6] That isn't a conceptual response. One cannot help but feel the relevant bodily disturbance except as good. One is hard-wired by nature to experience it in this way. It is not difficult to fathom why.

Pain is the opposite, as is the experience of smelling vomit, for example. In the latter case, one experiences a certain smell as bad, as something apt to harm. That's part of the representational content of the experience. In the former case, it is tissue damage within one's own body one experiences as bad, as apt to harm.

It is not just the affective dimension of pain that does not require concepts. Intuitively, one does not need concepts to have a pain, period. Given the right stimulus, one feels pain, whatever concepts one has in one's repertoire. This is not to say that one's cognitive assessment of a situation can never influence or affect one's experience of pain. Obviously, it can. The point is simply that the basic experience of pain requires no cognitive sophistication. Humans are hard-wired to experience pain when they undergo tissue damage, whatever they think or believe.

V. On the relationship of pain phenomenology and representational content

Suppose you have a throbbing pain in your finger. You can certainly have a pain without noticing it, as, for example, when you are distracted for a moment by something else, but if you do notice a pain—if you are introspectively aware of it,—then your attention goes to *wherever you feel the pain* (in this case, to your finger). Your attention does not go to where your experience is (that is, to your head, if your experience is a physical thing) or to nowhere at all. In attending to your pain, you are directly and immediately aware of certain qualities, which you experience as being localized in your finger. Among these is a quality you want very strongly to stop experiencing.

The same is true even if you are feeling a pain in a phantom finger. Still, you are directly aware of a quality you strongly dislike, a quality that you *experience* as being in a finger, even though the finger no longer exists. The first point to stress, then, is that the qualities of which we are directly aware in introspecting pain experiences are not qualities *of* the experiences (assuming that there is no massive error), but qualities of bodily disturbances in regions where the pains are felt to be, if they are qualities of anything at all.

Not everyone agrees with these points. Stephen Leeds, for example, writes:

---

6 The suggestion that the pleasingness of orgasms is part of their representational content is made in Tye 1995c. It is also the view taken by William Seager (forthcoming).

> A toothache does not seem to present the world to me in anything like the way that looking at my tooth does: rather, although it seems in some way to present me with properties, these are not presented as properties of my tooth (and still less as properties of my brain or nervous system)—rather they seem to be presented as properties of (or at) point *s* in a somatic field. (2002, p. 6)

I agree with Leeds that there is a big difference phenomenologically between my seeing the tooth (in a mirror, say) while I have a toothache and my experiencing the toothache. But this difference is one that poses no threat to the thesis that when we introspect pain experiences, the qualities of which we are directly aware are not qualities of the experiences; for the properties presented in the two cases are very different (just as they are in the cases of seeing a cube and feeling a cube by running one's fingers over it). In seeing the tooth, I am aware of its color and shape, and I am unaware of the qualities I dislike so much in having the pain experience. In experiencing the toothache, the situation is reversed.[7]

Leeds asserts that in experiencing a toothache, the properties presented to me are not presented as properties of my tooth. I disagree. When I focus upon my toothache, I am aware of an ache *in* my tooth, and in being aware of this ache, I am aware of a quality that I strongly dislike, whatever further account we give of its nature,—a quality that is presented to me *as* inside my tooth, as filling the region occupied by my tooth (the tooth to which I direct my dentist when he asks me "Where is the pain?").

Perhaps the thought is that I am introspectively aware of pain in my tooth, but I am aware of this by being aware of something else, a quality or qualities of a region of a somatic field. However, this seems to me no more plausible or well motivated than the corresponding claim that our visual experience of external physical surfaces is mediated by experience of visual sense-data.

In the case of vision, the surfaces of which I am directly aware are experienced by me at varying distances away and orientations. Unless there is some drastic error in visual experience, that is where such surfaces normally are. They are physical surfaces. They can be photographed. We ordinarily assume that others directly see those surfaces too.

In the case of pain, there is, I grant, a body image; and, in one way of talking, the pain in the finger is located on that part of the body image representing the finger. But all this really means is that there is, for each of us, a continuously updating sensory representation of the sort found in general bodily feeling, and that the experience of pain represents the quality or qualities felt as unpleasant as being instantiated at a certain location within the body space represented by the former representation.

Barry Maund (2003) has recently contested these claims. He comments:

> . . . the pain one feels in one's leg is a subjective feeling that one 'projects' onto the leg. It is not a real projection. One has a body image which represents the body. The pain is projected on to, is located on, that part of the

---

7 For more on the difference between the two cases, see Section VII.

> body image which represents the leg. Likewise with colors. There is a subjective quality which one 'projects' onto an external object, say to the moon, to represent it as yellow.

I confess that do not understand any of this. If I see the moon, I am not aware of a subjective visual field that represents the moon. I am aware of the moon and perhaps some stars located in distant regions of space before my eyes. Likewise, if I have a pain in my leg, I am not aware of an image that represents my leg. I'm aware of my leg and its condition. To suppose that it is the representation itself—the subjective visual field or the body image—of which I am really (directly) aware in these cases is like supposing that if I desire eternal life, what I really (directly) desire is the idea of eternal life. That, however, is *not* what I desire. The idea of eternal life I already have. What I desire is the real thing. And it does not help, of course, to say that it *must* be the representation of which I am aware, since the case might be one of hallucination—no moon or no leg—for patently, if there is no eternal life, it *still* isn't the idea of such a life that I really desire. If the pain is a phantom one or the visual experience totally delusive, I simply undergo an experience which represents something that isn't there.

It seems to me, then, that the right thing to say is that when I attend to a pain in my finger, I am directly aware of a certain quality or qualities as instantiated in my finger. Moreover, and relatedly, the only particulars of which I am then aware are my finger and things going on in it (for example, its bleeding). My awareness is of my finger and how *it* feels. The qualities I experience as bad or unpleasant are ones the finger or part of the finger or a temporary condition within the finger apparently have. My experience of pain is thus transparent to me. When I try to focus upon it, I 'see' right through it, as it were, to the entities it represents. But when I introspect, I am certainly aware of the phenomenal character of my pain experience. On the basis of introspection, I know what it is like for me phenomenally on the given occasion. Via introspection, I am directly aware of certain qualities which I experience as being qualities of my finger or episodes inside it, qualities which I experience as very unpleasant, *and thereby* I am aware of the phenomenal character of my experience. By being aware of these qualities, I am aware of what it is like for me. This is not to say, of course, that I *infer* the phenomenal character of my experience from my awareness of such qualities. Obviously, no *reasoning* is involved. Still, by attending to what is outside my head,[8] in my body, I know what it is like for me. So, my awareness of phenomenal character is not the direct awareness of a quality of my experience. Relatedly, the phenomenal character itself is not a quality of my experience to which I have direct access.

This conclusion could be resisted if one took the view that there is a massive error in pain experiences, that the qualities represented as instantiated in the body by such experiences do not belong there but are rather qualities of the experiences themselves. That is highly counter-intuitive, however, just as is the analogous position for color experience. And it is far from clear that it is intelligible. Take, for example, a sharp, dagger-like pain. The unpleasant quality one experiences is experienced as distributed

---

8 Ignoring headaches.

throughout, as filling, a dagger-shaped region of the body. Could that quality really be a quality of something that isn't extended at all? Is that conceivable? It seems to me conceivable that dualism is true and that my experiences are non-physical. But is it conceivable that the quality I experience as instantiated throughout a spatial region is really a quality of something that isn't spatial? Whatever we say to this, I know of no convincing arguments that pain experiences involve massive error, and I shall assume that the intuitive view is correct.

Now, as I experience the pain and I attend to it, necessarily, if any of the qualities of which I am *directly* aware changes, then the phenomenal character of my experience changes. If, for example, my finger starts to sting or to ache, the phenomenology of my pain experience changes. Why should this be? Facts like this one are surely not brute. Moreover, they obtain even in the case that I am undergoing a phantom finger pain. An explanation is needed of why the phenomenal character of pain experiences is sensitive in this way to qualities that, if they are qualities of anything at all, are qualities of bodily regions where the pains are experienced to be. The explanation surely is that the phenomenal character *involves* the bodily qualities of which the subject of the visual experience is directly aware, that these qualities at least partly *constitute* phenomenal character.

What, then, is the phenomenal character of pain? One possible hypothesis is that it is a subjective quality (or cluster of qualities) of the relevant bodily region, so that the experience itself inherits its phenomenal character from the relevant subjective quality (or qualities). That hypothesis will not accomodate phantom limb pain, however; for in that case there is no appropriate bodily region. The best hypothesis, I suggest, is that the phenomenal character of pain is representational content of a certain sort, content into which the experienced qualities enter. This explains why pain phenomenal character is not a quality of an experience to which we have direct access (representational content is not a quality of the thing that has representational content) and why pain phenomenal character necessarily changes with a change in the qualities of which one is directly aware (changing the qualities changes the content). It also explains why the phenomenal character of a pain experience is something the experience *has*, something that can be common to different token experiences, and why pain experiences have phenomenal character even if nothing really has the qualities of which one is directly aware via introspection.

VI. Phenomenal Content

The phenomenal character of a given pain, on my view, is one and the same as the phenomenal content it has. The phenomenal content of a pain is a species of representational content. It is content that is *nonconceptual*, not just in the sense that the subject of a pain experience, in undergoing it, need not possess any of the concepts that we, as theorists use, when we state the correctness conditions for that experience but in the full-blooded sense of being content that is different in kind from conceptual content.

Pain experiences represent in something like the way that the height of a mercury column in a thermometer represents. The latter normally tracks a certain temperature in the surrounding air, however that temperature is conceived. Similarly, pain experiences normally track tissue damage in regions of the body, however the tissue damage is conceived. The representational content pain experiences have via such tracking is not content that could be the content of a belief.

The phenomenal content is also *abstract*. That is to say, it is content into which no concrete things enter. Thus a pain in a leg cannot be an experience into which one of the subject's legs enters. The reason is straightforward. One can feel a pain in a leg, while lacking a leg, that is phenomenally indistinguishable from such a pain when the leg is present. The phenomenal character of the two experiences is the same, and thus, on my view, their phenomenal content. But in one of the two cases, there is no leg.

Another condition on the phenomenal content of pain is that it be suitably poised. In particular, it must stand ready and available to make a direct impact on relevant desires the subject has, for example, the desire to get away from the noxious stimulus. Moreover, the qualities in the content must be ones of which the subject of the experience is directly aware via introspection (assuming the subject is sophisticated enough to be able to introspect).

These three conditions on phenomenal content—that it be nonconceptual, abstract, and poised—constitute its essence (Tye 1995, 2000).

VII. Some remaining objections

One objection that has been raised to this proposal is that it cannot distinguish between the phenomenology of seeing one's damaged leg and feeling pain there.[9] This, I maintain, is not so. One's visual experience, as one views the leg, nonconceptually represents such features as color, shape, orientation of surface, presence of an edge. It does not nonconceptually represent tissue damage. One's pain does nonconceptually represent tissue damage, but it does not represent the other features.

I concede that, going by the phenomenal look of the leg, one will judge it to be damaged, and thus that it will look *to be* damaged or look *as if it is* damaged. This is a conceptual use of the term 'look,' however. The phenomenal sense of 'looks' is nonconceptual. It is captured by 'looks *F*,' where '*F*' is a term for a quality of which one is directly aware when one introspects the experience.

Maund (2002) has objected that the damaged leg *does* look damaged in the phenomenal sense of 'look.' He comments:

---

9 This is the Leeds' objection for the case of toothache, mentioned in Section I. It was first raised to me in conversation by Ned Block and John Searle; and it has recently been pressed by Barry Maund.

> . . . it seems to me that the phenomenal sense applies to a range of features or cluster of features besides the ones Tye cites: there are rusty looks, wooden looks, jarrah looks, metallic looks, and damaged-tissue looks.

I disagree. If one views something rusty in standard conditions, one is directly aware of a range of color and texture qualities, on the basis of which one judges that it is rusty. The object looks to one *to be* rusty; moreover it looks *like* other rusty things, and indeed, in one ordinary way of speaking, it looks rusty, but it doesn't look rusty in the nonconceptual, phenomenal sense. For rustiness isn't a quality of which one is directly aware when one introspects one's experience any more than is the quality of being feline, when something looks feline to one. Intuitively, felines and twin felines (molecule by molecule duplicates of cats belonging to a different species on a variant of Putnam's planet, twin earth) are phenomenally indistinguishable. In the phenomenal sense, they look alike. But cats look feline to us, whereas twin cats look twin-feline to the inhabitants of the other planet.

What is true here for cats is true *mutatis mutandis* for rusty things, wooden things, etc. And it is true for damaged tissue as well. Imagine in this case that on the twin planet, there is no tissue but an artificial look alike. Note, incidentally, that this is not to deny that "there is a characteristic experience-type that things which are rusty, wooden, jarrah, metallic and damaged tissue cause in optimal conditions" (Maund, 2002). There is such a type; but at the nonconceptual *phenomenal* level, that type is individuated by the cluster of qualities of which the subject is directly aware via introspection (qualities that are also represented by the experience). And those qualities are at the level of shape, color, texture etc. They do not include rustiness, woodenness, and so on.

A second objection concerns the concepts used in noticing one's own pains. While the experience of pain is nonconceptual, the awareness by the subject *that* he/she is experiencing pain requires concepts. Without concepts, those in pain would be 'blind' to their experiences. But which concepts exactly are the ones needed to notice pain? Murat Aydede (2001) replies:

> Certainly, not the concepts of bodily disturbance, tissue damage, disordered states, nociceptive stimulation, etc. All [most people (including small children)] need—and usually have in fact—is the concept of pain. But the concept of pain is not the concept of physical disturbance (tissue damage, etc) . . . (p. 16)

He continues, with the aim of raising a difficulty for my position:

> Consider a mild tickle, itch, tingle, ad pain occurring on the very same spot of my back for the same length of time (at different times, of course) . . . . They feel different. And if their feel is what they represent, then I seem to be able to discriminate between them, and single out their differences generally, without having the concepts that apply to what physically happens in the relevant part of my body, i.e. without having the concepts for what is represented by these experiences. The concepts I happen to have are none other then the concepts of an itch, tickle, tingle, and pain, and these con-

> cepts are not concepts of physical bodily happenings supposed to be represented by them (p. 16).

I agree with Aydede that 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. But why should this be thought to present a difficulty for my position?

I said earlier that when I experience pain, I am directly aware of the qualities the experience represents. By attending to these qualities, I am aware *that* I am experiencing pain. Introspection of experiences, in my view, 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. It is the reliability of this process that underwrites introspective knowledge of experiences.

Consider the case of introspection of thought contents, which is similar. If I think that water is wet, and I introspect, I become aware *that* I am thinking that water is wet. This awareness is not based upon an inference from other propositional states. Nor is it the result of attention to an internal auditory image of myself saying that water is wet, though such an image may accompany my thought. Intuitively, my introspective access to what I am thinking is direct. It seems plausible to suppose that introspection of thought contents is a reliable process that takes as input the content of the thought and delivers as output a belief or judgement that one is undergoing a state with that content (McLaughlin and Tye, 1998).

On this view of introspective knowledge of thought contents, the concept of a thought that *P* is, in its first-person present-tense application, a *recognitional* concept.[10] Those who have mastered the concept can introspectively recognize that an occurrent thought that *P* is present without going through any process of reasoning. In much the same way, we do not have introspective knowledge of our experiences by inferring it from something else. We acquire introspective knowledge that we are undergoing such and such an experience or feeling via a reliable process that triggers the application of a suitable concept. Experiential concepts—the concepts that enable us to form a first-person conception of our experiences via introspection—are, in my view, recognitional concepts of a special sort.

I agree with Aydede, then, that when I introspect a pain, the concept I apply is the concept PAIN or some more specific concept of that type, for example THROBBING PAIN. This concept may be purely phenomenal—and it will be, if all I notice is the phenomenal character of my state—but it may also be broader than that, applying to states that vary in their phenomenal 'feel.'

Since my awareness is that I am feeling pain and *not* that I am undergoing an experience that represents tissue damage, I do not apply the concept TISSUE DAMAGE via introspection. That tissue damage is the quality paradigmatically represented by pains qua sensory experiences is an empirical hypothesis. It is not a hypothesis, sup-

---

10 For more on recognitional concepts, see Brian Loar 1990.

portable by a priori reflection upon concepts or introspection. So, not only is the concept TISSUE DAMAGE not needed to feel pain, but it is also not needed to be introspectively aware that one is feeling pain.

There is an interesting question here, raised by Aydede, as to why, if the phenomenology of pains is given by their representational content (or some part of it), we have come to develop "an immediate epistemic and utilitarian/practical interest in the experiences in the first place rather than in the objects/conditions represented by them" (2001, p. 17). Aydede elaborates the concern further as follows:

> The explanandum is why our primary and immediate interest (epistemic as well as practical) is in the experiences themselves such that the related concept applies primarily to them rather than what they represent . . . we have no parallel of this phenomenon in standard exteroception, where the whole phenomenology can be accounted (assuming it can) for solely in terms of what is represented, i.e. the perceptual object and its qualities as the representational content of the relevant experiences (2001, p. 17).

What seems to be worrying Aydede is the thought that in the case of color experience, say, the information flow is such that red causes the experience of red, which then triggers the application of the concept RED to red, the color represented, but in the case of pain, the information flow goes from tissue damage to the experience of pain to the triggering of the concept PAIN, where that concept applies to the experience and not to the damage.

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.

Of course, if I see something red, I may judge that it is red, and here I apply the concept RED to the object seen. But likewise, if I cut my finger, I may judge that my finger is damaged, and in so doing apply the concept DAMAGE to my finger (or better its condition). The only difference between the two cases that I can discern is that when I introspect my color experience, the concept I exercise (EXPERIENCE OF RED) is a compound concept containing a concept for an entity represented by the experience, whereas in the case of pain, when I introspect, the concept I exercise (PAIN) has no such component.

A final objection I wish to consider is that the painfulness of pains has not been fully accounted for within the above theory.[11] Let me begin my reply by noting first that when

---

11 Aydede presses this charge in his 2001.

we say that something is painful, what we normally mean is that it *causes* the feeling of pain. Thus, we speak of a cut, a bruise, a sore, an operation, a cough, a lashing as painful, that is, as causing pain. This is why, if a person feels a pain in his left arm, and his doctor informs him that the cause of the pain lies in a disturbance in his heart, we allow that what is *really* painful here is the disturbance in my heart. If the doctor has made no mistake, *that* is what is actually hurting the patient.

It is evident that pains themselves are not painful in this sense. Pain does not cause the feeling of pain. So, it is no criticism of my view to say that the painfulness of pains, in this sense of painfulness, is unaccounted for within the representationalist theory.

What, then, has representationalism left out? If what is meant by the painfulness of pains is their phenomenal feel, then *that* hasn't been left out. For a theory has been offered of its nature, a theory that identifies the feel with a certain sort of sensory/affective representational content, the possession of which by pains explains why we normally have such a strong aversive cognitive reaction to them.[12]

So, representationalism seems to me to have much in its favor and no obvious problems. I continue to believe that it provides the best hope for a defensible, naturalistic theory of experience generally.

Michael Tye
The University of Texas at Austin

References

Armstrong, D. 1962 *Bodily Sensations*, London : Routledge and Kegan Paul.

Aydede, M. 2001 "Naturalism and Direct Realism about Pain," *Consciousness and Emotion*, 2, 1-24.

Block, N. 1983 "Mental Pictures and Cognitive Science," *Philosophical Review*, 92, 499-541.

Block, N. 1996 "Mental Paint and Mental Latex," *Philosophical Issues*, 7, E. Villenueva, ed., Northridge: Ridgeview Publishing Company.

Chalmers, D. 1996 *The Conscious Mind*, Oxford: Oxford University Press.

---

12 I think it is fair to say that in some of my previous writings, I have not emphasized enough the noncognitive, affective side of pain and its relevance to the overall content of pain experiences. But the importance to phenomenology of the affective dimension of content is clearly noted in Tye 1995c, for example.

Chapman, C. and Nakamura, Y. 1999 "A Passion of the Soul: An Introduction to Pain for Consciousness Researchers," *Consciousness and Cognition*, 8, 391-422.

Dennett, D. 1978 "Why You Can't Make a Computer that Feels Pain," in *Brainstorms*, Cambridge, Mass: The MIT Press.

Harman, G. 1990 "The Intrinsic Quality of Experience," in *Philosophical Perspectives*, 4, J. Tomberlin, ed., (Northridge: Ridgeview Publishing Company).

Jackson, F. 1977 *Perception*, Cambridge: Cambridge University Press.

Loar, B. 1990 "Phenomenal States," in *Philo sophical Perspectives*, 4, J. Tomberlin, ed., (Northridge : Ridgeview Publishing Company). Armstrong, D. Block, N. 1983

Leeds, S. 2002, "Perception, Transparency, and the Language of Thought," *Nous*, 00-00.

Maund, B. 2002 "Tye on Phenomenal Character and Color," web symposium on Tye: Consciousness, Color, and Content at: *http://host.uniroma3.it/progetti/kant/field/tyesymp.htm*.

Maund, B. 2003 "Tye: Consciousness, Color, and Content," *Philosophical Studies*, 249-260.

Melnyk, A. 1998 "Reply to Schroeder," presentation at APA, Pacific Meeting.

Melzack, R. and Casey, K. 1968 "Sensory, Motivational, and Central Control Determinants of Pain: A New Conceptual Model," in *The Skin Senses*, ed. by D. Keshalo, Springfield, Illinois.

Melzack, R. and Wall, P. 1983 *The Challenge of Pain*, New York: Basic Books.

McGinn, C. 1982 *The Character of Mind*, Oxford: Oxford University Press

McLaughlin, B. and Tye, M. 1998 "Is Content-Externalism Compatible With Privileged Access?" *Philosophical Review*, 107, 349-380.

Noordhof, P. 2001 "In Pain," *Analysis* 61, 95-97.

Noordhof, P. 2002, "More in Pain," Analysis 62, 153-154.

Price, D. 1999 *Psychological Mechanisms of Pain and Analgesia*, Progress in Pain Research and Management, Vol. 15, Seattle: IASP Press.

Quine, W. 1960 *Word and Object*, Cambridge, Mass: Harvard University Press.

Seager, W. forthcoming "Emotional Introspection."

Tye, M. 1995 *Ten Problems of Consciousness*, Cambridge, Mass: The MIT Press, Bradford Books.

Tye, M. 1995b "A Representational Theory of Pains and their Phenomenal Character" in *Philosophical Perspectives*, 9; also reprinted with revisions in *The Nature of Consciousness: Philosophical and Scientific Debates*, ed. by N. Block, O. Flanagan, and G. Guzeldere, Cambridge, Mass: The MIT Press, Bradford Books.

Tye, M. 1995c "Blindsight, Orgasm, and Representational Overlap," *Behavioral and Brain Sciences*,18, 268-69.

Tye, M. 2002, “On the Location of a Pain,” *Analysis* 62, 150-153.