

Is the Experience of Pain Transparent? Introspecting Phenomenal Qualities

Murat Aydede
Department of Philosophy
University of British Columbia

ABSTRACT. I distinguish between two claims of transparency of experiences. One claim is weaker and supported by phenomenological evidence. This I call the Transparency Datum. Introspection of standard perceptual experiences as well as bodily sensations is consistent with, indeed supported by, the Transparency Datum. I formulate a stronger transparency thesis that is entailed by (strong) representationalism about experiential phenomenology. I point out some empirical consequences of strong transparency in the context of representationalism. I argue that pain experiences, as well as some other similar experiences like itches, tickles, orgasms, hedonic valence, etc., are not transparent in this strong sense. Hence they constitute empirical counterexamples to representationalism. Given that representationalism is a general metaphysical doctrine about all experiential phenomenology for good reasons, I conclude that representationalism about phenomenal consciousness is false. Then, I outline a framework about how the introspection of phenomenal qualities works in light of the Transparency Datum, but consistent with the rejection of strong transparency. The result is a form of qualia realism that is naturalist and intentionalist (weak representationalist), and has close affinities to the adverbialist views developed in the latter part of the last century. I then apply this framework to pain experiences and their bodily locations.

Perceptualism about pain is the view that feeling pain in a body part is perceiving an extra-mental (physical) condition of that part. This view should not be confused with the mundane observation that feeling pain in a body part often conveys information about the physical condition of that part. This latter view is a platitude. The former is a substantive philosophical thesis that has only recently found an ever increasing number of defenders in the history of philosophy.¹ In this regard, it may be compared to what we might call ‘perceptualism about colour’ (indeed, about all so-called secondary qualities), namely, the view that seeing colour on a given surface is perceiving an extra-mental (physical) condition of that surface. Again, this should not be confused with the platitude that seeing the colour of a surface often conveys information about some physical condition of that

¹ See, for instance, Armstrong (1968), Pitcher (1970), Hill (2009).

surface. Unlike the latter, the former is a recent substantive philosophical thesis.² Although by historical standards, both kinds of perceptualism are quite radical and controversial, perceptualism about pain suffers from a *prima facie* implausibility that perceptualism about colours seems immune to. That is because pains are necessarily experienced or sensed:³ when it comes to pains, *esse is percipi*. But then a pain in a body part cannot, it seems, be an extra-mental condition of that part.

I have argued against perceptualism about pain elsewhere (Aydede 2009). Here I want to argue against a natural descendant of this view, *representationalism* about pain,⁴ according to which the phenomenal character of pain experiences is entirely determined by — or, one and the same thing as — the (wide) representational content of such experiences. Thus, pain experiences represent (say, non-conceptually) certain kinds of bodily events or conditions. The phenomenal character of a pain experience, then, is its representational content — that a certain kind of bodily disturbance is occurring in a part of one's body.⁵ Again, this representationalist view should not be confused with the truism that our pain experiences often convey information about the physical conditions of our body parts.

Representationalism about pain is the result of a more general representationalist approach to experiences, according to which the phenomenal character of *all* experiences can be reduced to their (wide) representational content. With the assumption that the kind of representational content involved in experiences can be naturalistically accounted for, representationalism in philosophy of perception is a metaphysical project aiming to provide a reductive explanation of phenomenal consciousness. On this view, experiential

² See, among others, Armstrong (1968), Smart (1975), Hilbert (1987), Byrne & Hilbert (1997, 2003).

³ Consciously or unconsciously. If it is possible to have unconscious sensations, then pains may be unconscious but still they are necessarily sensed. This is not something true of colours and other secondary qualities: there is not even a remotely plausible sense in which colours themselves can be unconscious, although a sensation of colour can be — if sensations can be unconscious. So I will leave this issue aside, and for convenience, assume that all sensations/experiences are conscious.

⁴ Although representationalism is a natural descendant of perceptualism, neither view, strictly speaking, implies the other. So the issues tend to be somewhat different as representationalism is a metaphysical thesis about all phenomenal character — see below. Nevertheless, defenders of these views are natural allies.

⁵ Some representationalists (e.g., Bain 2013; Cutter & Tye 2011) additionally claim that the affective (awful, unpleasant, painful) phenomenology of pain experiences is also representational and that the content is something like: [that bodily disturbance is bad for one].

phenomenality is nothing but the naturalistically kosher representational content of experiences.

The plan for the paper is as follows. In the next section, I further examine varieties of representationalism and clarify the version that will be the target of my criticism, and spell out what it says of pain that makes the view so radical. This is important because of two things: one is the need to distinguish weaker forms of representationalism free of the difficulties I will raise for the stronger version. The second is to analyze the commitments of the stronger representationalism to see some of its consequences that will play a key role in my argumentative strategy.

Section 2 is meant to be a contribution to the ongoing debate about the so-called perceptual transparency. I distinguish between two claims of transparency of experiences. One claim is weaker and supported by phenomenological evidence. This I call the Transparency Datum. Introspection of standard perceptual experiences as well as bodily sensations is consistent with, indeed supports, this Datum. I then formulate a stronger transparency thesis. I will show that strong representationalism entails this stronger transparency thesis, but is not supported by the Transparency Datum over its competitors.

Section 3 is critical in drawing out two empirical consequences of strong representationalism — given its entailment of strong transparency and the view it is committed about introspection. The subsequent two sections (§§ 4–5) will then show how pains as well as some other experiences constitute empirical counterexamples to strong representationalism because of these consequences. As far as I am aware, this is a new form of argument against strong representationalism that has not been made in the literature before.

The Transparency Datum is widely accepted in the literature on perception as it is strongly supported by phenomenological evidence. Any theory of perceptual experience must explain, or at least accommodate, this Datum. Strong representationalism with its entailment of strong transparency is *prima facie* an attractive way of explaining the Transparency Datum. So it is incumbent on those who reject strong representationalism to explain this Datum given the puzzles pain experiences pose. It is the job of Section 6 to do this, where I develop a positive account of perceptual experience along with an account of

its introspection which is broadly naturalistic, weakly representationalist, and is a version of qualia friendly adverbialist views developed in the second part of the last century. In subsection 6.1, I develop the general framework for perceptual experience at large, and in subsection 6.2, I apply the framework to pains and show how it can insightfully explain many puzzling features of pains and other similar experiences. Section 7 concludes the paper by drawing some general lessons.

1 Representationalism and Intentionalism

Strong representationalism (henceforth, representationalism — see fn. 6) is a stronger view than what is sometimes known in the literature as *intentionalism*, which is merely the view that experiences represent.⁶ Unlike representationalism, this more general view need not be reductive or naturalistic; it need not involve commitment to only wide or Russellian content — it could be narrow or Fregean. It is also typically more restricted in its scope than representationalism: it may allow for there being aspects of phenomenology that are not intentional at all (historically, typical examples include the phenomenal character of pains and other intransitive bodily sensations such as itches, tickles, and orgasms, as well as aspects of emotions and moods).

According to representationalism, however, the phenomenal character of *any* type of experience, indeed experiential phenomenality in all of its manifestations, is one and the same as (or, is entirely determined by) what is represented by these experiences and what is represented are always extra-mental objects/features of one's perceptual environment.⁷ If I am seeing a red round tomato, the roundness and redness are among the metaphysical determinants of the phenomenal character of my visual experience. Obviously, any visual experience of a red and round tomato has a lot more to its phenomenal character than is

⁶ Terminology here is not completely settled. The label 'intentionalism' is sometimes used interchangeably with 'representationalism' in the literature. Strong and weak representationalisms would then be mutually exclusive subspecies of representationalism or intentionalism. Many defenders of strong representationalism don't bother to mark their version as strong representationalist and use 'representationalism' to characterize their view. I will follow this practice and use 'representationalism' to mean strong representationalism. I will use 'intentionalism' to denote the more general view. When it matters to mark a position that is intentionalist but not strong representationalist, I will use 'weak representationalism'.

⁷ Among its defenders are Harman (1990), Dretske (1981, 1995), Tye (1995, 2000, 2006a), Byrne and Hilbert (1998, 2003), Jackson (2004), Byrne and Tye (2006). For a detailed development and defense of representationalism about pain and pain affect in particular, see Tye (1995, 1996a, 1997, 2006a, 2006b), Bain (2003, 2007, 2013), O'Sullivan and Schroer (2012).

contributed by the redness and roundness of the tomato. Whatever more there is, *all* of it is *entirely* determined by extra-mental features of the scene (containing the tomato) that are (widely) represented by my experience. Representationalism is thus meant to rule out two claims:

- (a) there can be non-representational (elements of) phenomenology, and
- (b) whatever extra-mental objects/properties an experience with a given phenomenology represents, it represents them contingently.⁸

Representationalism thus implies phenomenal externalism, the claim that phenomenal character is determined *solely* by factors that are completely external to one's experiences or one's mind. Intentionalism, by contrast, is compatible with either or both of these claims. There are forms of intentionalism that are internalist about phenomenal character — phenomenal internalism is the denial of phenomenal externalism.⁹ For instance, if qualia are intrinsic and directly introspectable phenomenal qualities of experiences, then various naturalist or non-naturalist qualia views are compatible with intentionalism, but not with representationalism.

Given that phenomenal character is essentially introspectable for creatures like us with relevant mental capacities, and according to representationalism, is entirely metaphysically determined by (wide) representational content, introspection of experiences cannot reveal any phenomenal quality as belonging to the experiences as opposed to belonging to the extra-mental world. For if introspection were to reveal any phenomenology-determining quality that could not be entirely attributed to those elements of the extra-mental world represented by experiences, then *something other than wide*

⁸ In Ned Block's terms, (a) amounts to the existence of mental latex, and (b) to the existence of mental paint (Block 1996).

⁹ Thus, according to my usage, those positions that take external factors to be (merely) contributing factors to the determination of phenomenology count as phenomenal internalist. Phenomenal internalists need not deny that external factors are among the determinants of experiential phenomenology. One need not be a representationalist to endorse phenomenal externalism: disjunctivism and naive realism about perception, as well as versions of behaviorism about perceptual states, are forms of phenomenal externalism that are not (typically) representationalist. Thus, although my main target in this paper is representationalism, I take the main line of argument based on features of pain experiences to be equally effective against disjunctivists: as long as they are not eliminativist about perceptual phenomenology, they are committed to phenomenal externalism, and with it, to a strong form of transparency.

representational content would be among the phenomenology-determining elements of experiences. But this would refute representationalism. Thus representationalism implies phenomenal externalism, and with it, a strong form of *transparency* of experiences. If this form of transparency is false, then representationalism is false.

Understood in this way, representationalism has consequences that make *representationalism* about pain even more radical and controversial than *perceptualism* about pain. If representationalism is true, then experiences are transparent in a very strong sense: introspecting them does not involve any direct or immediate access to any sensation or sensation-like phenomenal quality of an experience. Applied to pain, this transparency is well expressed by Tye:

... 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 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. ... So, my awareness of [pain's] 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. (Tye 2006a: 110)

This, combined with representationalism, implies a straightforward appearance/reality distinction applied to pain: pain experiences are just like other perceptual experiences that (re)present the extra-mental reality around us (including the physical conditions of our bodies). Thus the qualities attributed to body parts in pain are physical qualities (bodily disturbance, tissue damage, etc.) that can be misattributed. Moreover, given representationalism, when I am feeling pain, say, in my finger, although I may be aware *that* I am having a pain experience, I am not aware *of* a sensation or an experience: what I

am directly aware *of* is an entirely physical condition of my body (whether or not this awareness is veridical).¹⁰

As far as I know, this sort of view about pain has never been seriously entertained by anybody throughout the long history of philosophy. Indeed, until about 30 years ago, claiming that pain experiences are transparent in this strong sense would have been considered *absurd*. To many (maybe even most) ears, it still sounds absurd today. In what follows I will argue that pain experiences are not transparent and therefore they are not representational. If my conclusions are right, then representationalism about experience as a metaphysical project needs to be abandoned since pains (and other similar phenomenal states — see below) are *counterexamples* to the general representationalist claim about experiential phenomenality.

Despite much recent discussion in the literature, however, the alleged transparency of experiences seems not a well-understood phenomenon.¹¹ It is supposed to be the kind of phenomenon that is evident on the basis of careful, reflective introspection, not a controversial philosophical thesis. Thus the transparency of experiences is supposed to be an empirical datum, revealed by introspection, that needs to be explained or at least accommodated by theories of phenomenal character. Most representationalists take this datum to argue against phenomenal internalism. There are, however, defenders of internalism (intentionalist or otherwise) who think that the transparency of experiences as a datum does not pose any serious challenge to their position. Given that phenomenal internalism has been historically the dominant view and is still very popular, it's puzzling that there has been no serious attempt to explain how the apparent transparency of perceptual experiences can be squared with any version of internalism. So we need to have a better understanding of transparency.

¹⁰ I am following Dretske (1999) and Tye (2002) in distinguishing between *awareness-of* and *awareness-that*. The latter awareness, unlike the former, requires concepts. Although the focus of this paper is representationalism about pain, the issues raised by pain are not peculiar to pain. As already mentioned, there are various other sensations (other 'intransitive' bodily sensations such as aches, itches, tickles, tingles, orgasms, dizziness, etc.) and phenomenal occurrences (moods, emotions, and sensory affect such as pleasantness or unpleasantness or even painfulness of certain sensations) that raise exactly the same difficulties for representationalism. I will come back to this below in §5.

¹¹ Despite the existence of quite helpful literature on transparency (Crane 2000, Tye 2002, Kind 2003, Siewert 2004, Stoljar 2004, Hellie 2006, Nida-Rümelin 2007), there is, it seems to me, still a lot of disagreement about what exactly the phenomenon is and what its significance is for theories of perception.

2 Transparency Datum and Representationalism

Let us start with some agreements among the defenders of both phenomenal externalism and internalism. Both Thomas Reid and G. E. Moore were (or would have been) on the externalist side of the divide. Yet, both point out that (some) experiences are transparent or, as Moore famously put it, *diaphanous*. Here is Thomas Reid:

We are so accustomed to use the sensation [of hardness] as a sign, and to pass immediately to the hardness signified, that, as far as appears, it was never made an object of thought, either by the vulgar or by philosophers ... There is no sensation more distinct, or more frequent; yet it is never attended to, but passes through the mind instantaneously, and serves only to introduce that quality in bodies, which, by a law of our constitution, it suggests. (Reid, *An Inquiry into the Human Mind*, 1764/1872: 120)

Here is G. E. Moore:

[T]hough philosophers have recognized that something distinct is meant by consciousness, they have never yet had a clear conception of what that something is ... [T]he moment we try to fix our attention upon consciousness and to see what, distinctly, it is, it seems to vanish: it seems as if we had before us a mere emptiness. When we try to introspect the sensation of blue, all we can see is the blue: the other element is as if it were diaphanous. (Moore 1903: 450)

Gilbert Harman and Michael Tye, both defenders of representationalism and strong transparency, eagerly agree and expand. Thus Harman says:

When Eloise sees a tree before her, the colors she experiences are all experienced as features of the tree and its surroundings. None of them are experienced as intrinsic features of her experience. Nor does she experience any features of anything as intrinsic features of her experiences. And that is true of you too ... Look at a tree

and try to turn your attention to intrinsic features of your visual experience. I predict that you will find the only features there to turn your attention to will be features of the tree. (Harman 1990: 667)

Here is Tye:

Focus on some object that you recognize, a blue disk, say. Now turn your attention inwards and try to pick out intrinsic features of your experience, inside you, over and above what it is an experience of. Try to discern intrinsic features of the mental paint. The task seems to me impossible. In turning one's attention inwards, one seems inevitably to end up focusing on external features one's experience represents the object as having, to the blueness and roundness as out there in the world covering or framing the surface of the object. In this way, the experience seems diaphanous or transparent. The same point holds, even if you are hallucinating. (Tye 1996: 295–296)

Momentarily leaving aside the issue of the scope of the transparency of experiences (whether it holds for all aspects of experiences of any phenomenal kinds or whether it holds for only some — say, only of clearly *perceptual* experiences excluding perhaps sensory affect, emotions and moods, etc.), they all seem to agree about two points: one about the apparent location of qualities and the other about the direction of focus or attention.

(LOCATION) The qualities that we are aware of in virtue of having a (perceptual) experience (for instance, in seeing a tree or a blue object, or feeling the hardness of an object) all *appear* to be qualities of extra-mental objects (particulars), including bodily parts.

This is supported by introspection. It describes *how* your (perceptual) experiences present the qualities they do: they present them as qualifying extra-mental *particulars/objects* — or

perhaps better, they represent them as instantiated by particulars external to one's experiences or mind.¹²

(FOCUS) If there are intrinsic qualities *of* experiences, it seems impossible to attend to or focus on these qualities *without* attending to or focusing on the qualities that these experiences present *as* belonging to the extra-mental particulars.

Conditionalized in order to find a common denominator between the defenders of phenomenal internalism and externalism, FOCUS is a claim about our own epistemic/psychological *capacities* that needs to be interpreted as implying something like: the introspective evidence is consistent with the claim that it is impossible to *directly* attend to the intrinsic qualities (if there are any) of our experiences.

The agreement among these four authors would come to an end, probably, if we consider whether LOCATION and FOCUS can be maintained with a maximally general answer to the scope question. Representationalists like Harman, Dretske, and Tye are bound to make these claims regarding *any* kind of experience with a phenomenology (or *any* kind of phenomenology an experience can have). But it is doubtful whether including, for instance, affective phenomenology (such as the pleasantness or unpleasantness *of* experiences) in the scope of LOCATION and FOCUS would make these claims still supportable on the basis of introspection alone. Nevertheless, I will leave the scope question aside in what follows and assume that LOCATION and FOCUS are supported by introspection even under the most general answer to the scope question. This is not because I believe that LOCATION and FOCUS would remain true under the widest scope, but because my argumentative strategy will rely on different considerations.

With this proviso, call the conjunction of LOCATION and FOCUS, the Transparency Datum — the Datum, for short. Note that the Datum, all by itself, does not make any claim about the metaphysical status of the *qualities* said to be attributed to the external particulars or to the experiences (if there are any so attributed), and in particular, about whether the

¹² Although more to be said about the intended meaning of 'appear' in LOCATION, the intended meaning, unlike in FOCUS that follows ('seems'), is phenomenological (not epistemological). Also, here and in what follows, I will put aside versions of idealism (if there are any) that would not allow for the existence of non-mental particulars.

qualities themselves are objective or non-mental (how is the introspection itself supposed to disclose something like that?).¹³ Also note that there is a certain element of triviality to both LOCATION and FOCUS. Regarding the former: if we have in mind *perceptual* experiences, it is quite natural to expect that *perception* would attribute the qualities it registers to the extra-mental particulars and that *introspection* would reveal this to be so about *perception*. Regarding the latter: it is also *prima facie* natural to expect that whatever the sense in which we can be said to directly attend to or focus on the qualities experientially presented to us as qualifying extra-mental particulars (in sensing or perceiving them), it is not quite in the same sense in which we can, in introspection, attend to or focus on the intrinsic qualities of our experiences (if there are any).

So, when read carefully, I believe that the Datum can be established on the basis of introspective evidence. Indeed, I think it is clearly true (at least for clearly perceptual experiences). However, this Datum is too weak to support representationalism or phenomenal externalism over its internalist competitors. For instance, after noting the transparency of experiences, perhaps in a way that would align them with the Datum, Reid and Moore make remarks that indicate that they would reject any stronger transparency thesis of the sort the representationalists are committed to. Right after the passage quoted above, Reid writes:

Sometimes it is easy enough to attend to the sensation occasioned by the hardness of a body — for instance when it is so violent as to occasion considerable pain. In that case nature calls our attention to it, and we then acknowledge that it is a mere sensation that can't exist except in a sentient being. (Reid 1764/1872: 120)

Similarly, it is well known what Moore says right after the passage quoted:

¹³ However, as a matter of fact, all representationalists believe that the *qualities* attributed to extra-mental particulars by experiences are objective and non-mental. So, for obvious reasons, among strong representationalists we don't find, for instance, defenders of subjectivist, dispositionalist or relationalist views of secondary qualities. I will also assume, again along with all representationalists, that the *particulars* sensed or perceived are completely objective and non-mental (physical). I am aware that sense-datum theories raise delicate issues about transparency, but I cannot address them here.

Yet it [the sensation of blue] can be distinguished if we look attentively enough, and if we know that there is something to look for. (Moore 1903: 450)

Although there are interpretative issues about their overall position, both Reid and Moore, along with pretty much everyone else in their own times, thought that we can be directly aware of sensations apart from the features that these sensations are sensations of. Unless Reid and Moore, along with countless other respectable philosophers until the present day, are confused in an elementary sort of way,¹⁴ The Transparency Datum is quite compatible with a robust rejection of representationalism or phenomenal externalism — below in §6, I will present my own account of how.

What is the strong transparency thesis that representationalism implies? Alas, there are no clean formulations of this thesis in the writings of representationalists that would clearly distinguish it from the Datum. However, right after his last quotation above, Tye makes a claim that is as good as it gets:

When you introspect, you still seem to find yourself attending to external features, to what it is that your experience is of, even if, in reality, there is nothing before you. Generalizing, introspection of your perceptual experiences seems to reveal only aspects of what you experience, further aspects of the scenes, as represented. Why? The answer, I suggest, is that your perceptual experiences have no introspectable features over and above those implicated in their representational contents. So, the phenomenal character of such experiences — itself something that is introspectively accessible, assuming the appropriate concepts are possessed and there is no cognitive malfunction — is itself representational. (Tye 1996a: 296)

So let us formulate the Strong Transparency claim to be:

(S-TRANSPARENCY) Experiences have no introspectable features over and above those implicated in their representational contents.

¹⁴ From the writings of some representationalists, sometimes one gets the impression that they really do think these philosophers were indeed so confused.

We need to have a clearer understanding of what this means. Evidently, according to representationalists, there *are* introspectable features of experiences that are *not* over and above those implicated in their representational content. What are these? Suppose I am looking at a red and round tomato in good light. Let us suppose that I am having a veridical visual experience of the tomato as having a determinate redness and roundness. My experience, *e*, represents the tomato as red and round. Of course, *e* represents many other visually detectable determinate features of the tomato and its surround. Let us concentrate on redness and roundness. Redness and roundness are part of the (non-conceptual) representational content of my experience. So clearly they are ‘implicated’ in *e*’s representational content. In what sense do I introspect them? I introspect them *as features of the tomato (or its surface) that are represented in my experience* (or, as features being experienced by me), not as instantiated in my experience — just as LOCATION says. It is the tomato that instantiates them and I introspect *that* my experience (re)presents these qualities as instantiated by the tomato.¹⁵ Redness and roundness are features of the tomato, not of my experience.

Furthermore, according to representationalists, my introspection of this intentional feature of my experience (being an experience as of red and round) is not a result of my direct attention to my experience. In other words, it is not a result of my attending to my experience itself *without* attending to the redness and roundness of the tomato. I attend to *what* my experience represents and somehow come to have the introspective knowledge *that* I am having an experience as of red and round.¹⁶ But this is consistent with, indeed encouraged by, FOCUS.

But what are the introspectable features whose existence is denied by Tye (in S-TRANSPARENCY) and affirmed by the likes of Reid and Moore? How can we settle this

¹⁵ I will leave aside whether experiences can represent such high level properties as being a tomato or a pine tree. The controversial issues surrounding transparency concern low level (usually sensorially detectable) properties and their representation. Also, almost all representationalists consider the representational content of experiences to be non-conceptual — whatever exactly this comes to. This is a point of agreement between me and representationalists, and what I say below about what is required for introspection does not contravene this.

¹⁶ This view of introspection is sometimes known (due to Dretske 1995) as the Displaced Perception Model (DPM) of introspection. For a critical discussion, see my (2002) and Aydede & Fulkerson (2014).

question: of any alleged introspectable feature of an experience, what is it to be over and above those implicated in the representational content of that experience? It seems that, once we grant that the Datum constrains experiential phenomenology and its introspection, it is practically impossible to find evidence for the claim that there are some (phenomenology-determining) features of some experiences that are both introspectable and *not* implicated in the representational content of those experiences (i.e., features that are not represented by experiences but instantiated by them).

However, concluding that there is no such evidence would be a mistake. Representationalism and S-TRANSPARENCY have empirically testable consequences. We need to carefully describe what these consequences are and see whether evidence bears them out. I now turn to this task.

3 Representationalism and the Introspection of Experiences

Representationalism about experiences, given its internal commitments, has consequences for the *introspection* of these experiences. As we have seen, if the phenomenal character of experiences is entirely determined by their (wide) representational content and the phenomenal character is introspectable, then introspective knowledge of experiences is exhausted by a specification of their representational content. There is no other way. This specification may be (practically always) partial. So in the above example:

(P-CONTENT) *that* is red and round,

where ‘that’ refers to the tomato, is a partial specification of the representational content of my perceptual experience when I see the tomato. But, of course, P-CONTENT is not about my experience, it expresses a *de re* proposition about the tomato. If I come to *believe* it, say, as a direct response to my experience, my belief would be a *perceptual*, not introspective, belief. Thus, the content of my *introspective judgment* about my experience must be something like:

(I-CONTENT) I am experiencing that as red and round,

where ‘that’ refers to the tomato. Thus, according to representationalism, the capacity for making *introspective* judgments about one’s experiences requires the capacity to make *perceptual* judgments directly prompted by those experiences. And both of these capacities require possession of concepts. Without concepts one does not have the capacity to make judgments or form beliefs at all.¹⁷

More specifically, according to representationalism, *introspection* requires the possession of concepts needed to express what is represented by one’s experiences.¹⁸ The possession of these concepts by those capable of introspecting their experiences requires the capacity to make direct perceptual judgments with these concepts: without having a capacity to form *de re* judgments attributing perceptible properties to the particulars so represented in one’s experiences, one cannot come to possess concepts required to articulate what is perceived or introspected. For instance, those capable of introspection cannot come to possess the concept of red (RED)¹⁹ or the concept of round (ROUND) without having the cognitive *capacity* to make *de re* judgments such as ‘*this* is red’ or ‘*that* is round’ in direct response to one’s experience whose partial content is also [this is red] or [that is round], where the demonstratives pick out extra-mental particulars. This is an empirical consequence about *experientially* (or, *sensorially*) acquired concepts that express properties that can be directly represented in experiences.²⁰ But note what follows from this.

¹⁷ Again, this does not contravene the claim that the representational content of perceptual experiences is not conceptual. I agree with representationalists that one need not possess concepts in order to have sensory or perceptual experiences. This is, of course, consistent with our attempt to partially specify the non-conceptual perceptual content of experiences propositionally as in (P-CONTENT). As far as I can tell, all representationalists agree with the claim made in the main text — see below.

¹⁸ For the clearest and emphatic statement of representationalism’s commitment to the availability of perceptual concepts for introspection and its general rationale, see Dretske (1995: 138–140) and (1999: 18–20).

¹⁹ Following standard practice, I will capitalize the name of concepts, where concepts are understood to be species of mental representations in more or less the psychologists’ sense. These representations along with the sensory representations underlying perceptual experiences are presumed to be realized in or implemented by the relevant hardware of the central nervous system.

²⁰ True at least for the concepts of low level perceptual properties such as being red or being round that are uncontroversially representable in our visual experiences. It is certainly true for the concepts of so-called secondary qualities represented in the experiences that are generated by their relevant sensory modalities. Indeed, I am assuming that every sensory modality consciously interfaced with conceptual systems comes with a proprietary range of phenomenal qualities whose concepts would require the sensory modality in question for their acquisition and direct application. The acquisition of amodal concepts for high-level

According to representationalism, every experience attributes objective properties or features to extra-mental particulars. This is what generates their representational content that can be veridical or not. So for instance, if the partial content of my experience when I see the tomato can be expressed by P-CONTENT, then my experience is veridical only if the tomato is red and round. Similarly, with my *de re perceptual judgment* directly (and appropriately) prompted by my experience: it is true (in that context) if, and only if, the tomato is red and round. The veridicality conditions of my introspective judgment with I-CONTENT are quite different: it is true (in that context) if, and only if, I am having an experience with the (partial) content expressed in P-CONTENT. It does not matter whether P-CONTENT itself is true or false. My introspective judgment, in other words, tracks my experience, not what my experience represents, whereas my experience and my perceptual judgment based on it track whether the tomato itself is red and round.

This pattern of interplay among experiences, perceptual and introspective judgments is exactly what is predicted by representationalism and the introspection model it implies. More specifically, given representationalism, S-TRANSPARENCY has the following empirical prediction:

(P1) For any introspectable feature one can epistemically encounter in one's experience, and for anyone who is capable of introspecting that feature, one has a concept that has *de re labeling* uses in the sense that the introspecting subject can form *de re* perceptual judgments in which the concept is directly applied to whatever extra-mental particular one's experience represents as having the feature.

According to representationalists, such features as one can encounter in one's introspection are features represented in one's experience as being *extra-mental (objective)* features or conditions of particulars given in one's experience. This yields a second prediction about the *de re* perceptual judgments mentioned in (P1):

properties such as being a tomato or being a pine tree, or damage (see below) may also require the actual or potential *ability* to make *de re* judgments if certain forms of an informational psychosemantics are true, but I will leave this issue aside as these are not uncontroversially experiential concepts. See Aydede & Güzeldere (2005) for more discussion.

(P2) These *judgments* in which the concept for the feature is used to attribute the feature to a particular are true or false according to whether the particular has that feature or not.

P2 is in line with the parallel claim that the *experiences* that prompt such judgments are veridical or not according to whether the particulars they represent as having certain features do indeed have these features or not.

Now suppose a phenomenal internalist makes the following quite natural claim: for any introspectable feature of an experience of an introspection-capable subject, if it is such that its concept does not have *de re* labeling uses, then this feature is a feature of the experience that is over and above those implicated in the representational content of that experience. Why would this be natural? If we introspect a feature for which we do not have a *de re* labeling concept (in the sense specified in P1), we cannot (correctly or incorrectly) attribute this feature to an extra-mental object. But if we cannot, and if we seem to be applying the concept correctly to something nevertheless, this feature must be (at least, partly) a feature of the experience itself. But this would refute S-TRANSPARENCY, and thus representationalism that implies it. The question of whether or not such features exist, then, is the question of whether the predictions expressed by (P1) and (P2) are empirically borne out. To show they are not, all one needs to show is that there are experiences whose introspective report does not rely on concepts that have *de re* labeling uses. This, in effect, amounts to showing that there are experiences whose introspection reports features for which there is no appearance/reality distinction. Are there such experiences? Plenty!

4 Pain Experiences

I will get to other examples in a moment. But let us start with pain experiences. Representationalists like Tye think that pain experiences are transparent in that introspection of them does not reveal any quality over and above those implicated in their representational content. If I am feeling a pain in my finger and quite aware of it, I am attending (somatosensorially, rather than, say, visually) to my finger and to a certain quality it has. This quality is attributed by my experience to the finger and I am introspectively

reporting my experience as an experience *of* my finger's having this quality. Since I am attributing *pain* to my finger, *prima facie* this quality is pain. But if representationalism is true, this quality *is* in fact some physical (extra-mental) condition of my finger.

Representationalists typically identify this condition to be some sort of bodily disturbance or tissue damage, a completely objective property. According to representationalists, I may not know the exact nature of this objective quality except in a vague and most general sort of way as some kind of bodily disturbance distinct from tickles or tingles and somewhat similar to itches, etc. — just as I may not know what complex physical property red is except that it is some physical condition of surfaces distinct from that colour or this colour and similar to or different than this one or that one. It may be that science will discover the exact nature of these disturbances occurring in bodily locations — just as science has told us that red is (let us assume) a set of surface reflectances of such and such kind. The concept of pain applied to bodily parts may be just like the concept of red applied to physical surfaces in *de re* perceptual judgments of the form:

that is red,

where 'that' refers to an object or surface. If so, we should expect pain-attributing *de re* judgments to be false in cases where the bodily part does not have the physical quality attributed, i.e., where the bodily part is not in any way physically disturbed. Any genuinely *de re* perceptual judgment directly prompted by a relevant experience is false in case the property attributed both by the experience and *de re* judgment is not possessed by the particular that the property is attributed to.²¹

So, given representationalism, this is how one would expect the empirical world to *be* regarding pain experiences undergone by introspection-capable subjects and their judgments. Representationalism, although a philosophical thesis, is not a conceptual claim unconstrained by empirical facts. It concerns, among other things, the psychology of

²¹ It is interesting to note that we do not use locutions such as 'this is pain' or 'that is an itch' paralleling the *de re* perceptual judgments like 'that is red' or 'this is an apple'. 'This hurts' is a different matter whose discussion is complicated and requires more space than I have here — but see the second proposal in §6.2 below which comports particularly well with the use of this expression. Although I disagree with his final analysis, Bain (2007) has a useful discussion of this expression.

certain kinds of organisms with intentional capacities of certain complexity. In very broad brushes, it tells us, among other things, how the basic psychological capacities are organized and related to each other as a consequence of its claims about the nature of the psychological states generated when these capacities are exercised. So we can ask: is the world like what representationalism predicts? To answer this question, we need to carefully look at those who are capable of experiencing pain and of making certain kinds of perceptual and introspective judgments on the basis of their experiences, and see whether the way they are related is indeed the way predicted by representationalism. So what are the facts?

Well, facts are well-known. No pain-attributing judgments, if they are made appropriately on the basis of the relevant experiences, are false in virtue of the fact that there is no physical disturbance of any kind in the location where pain is attributed. People (folk and the scientists alike) who feel pain in a bodily part do, as expected, typically make immediate judgments based on their experiences attributing *pain* to those body parts. When made appropriately, these judgments are never taken to be false merely in virtue of the fact that there is no bodily disturbance of any kind in the bodily location to which pain is attributed.²² So consider, for instance, someone with a heart condition who feels pain in his left arm due to his relevant heart muscles not getting enough oxygen (a standard referred pain case). There is nothing physically wrong with the arm. When this patient judges and reports that there is pain in his left arm, no body with the typical mastery of the relevant concepts (pretty much every normal adult) would take this judgment to be false. In fact, come to think of it, if a doctor, knowing the relevant facts, were to judge that the patient's perfectly lucid and sincere claim is in fact false (because there is nothing physically wrong with his arm) and proceed accordingly (send him home — there is no pain in his arm), he or she might be sued for malpractice. I will not belabor this point any further, since it is not much in dispute (even by representationalists — see below). So,

²² Consider the empirical facts that linguists rely on when constructing and testing theories about the deep syntactic structure of natural languages. Some of these are facts revealed by ordinary speakers' actual grammaticality judgments. Similarly, the empirical facts that I claim falsify representationalism are facts revealed by people's (including scientists' and clinicians') judgments about pain. These reveal, I claim, the actual cognitive architecture of how pain experiences interface with conceptually structured cognition in people with the relevant sort of sensory and conceptual competency. This architecture is not one predicted by representationalism — on the contrary.

simply put, these pain-attributing judgments are, as a matter of fact, *not* taken to be true or false in virtue of the presence or absence of any physical disturbance,²³ or whatever is the objective condition of body parts that the representationalist wants to claim as the representational content of the pain experiences.

The way these pain-attributing judgments actually work is exactly the way introspective judgments generally work: they track experiences, not what these experiences represent if they represent a physical condition of body parts. In other words, the empirical facts about the interplay between pain experiences and the judgments they directly prompt is opposite of what representationalism predicts. Our actual pain-attributing judgments *are* already *introspective* judgments if there are no corresponding *de re perceptual* judgments attributing pain (= bodily disturbance, according to representationalism) to body parts whose correctness conditions track whether such disturbances are occurring or not. If pain experiences were strongly transparent in the sense required by representationalism, this would not happen — in fact it should not happen. But there they are: the world turns out not to comply. So (P1) and (P2) are false. Hence pain experiences are not strongly transparent.²⁴

Let me briefly comment on a couple of ways a representationalist might respond. Given how well known the facts are about pain-attributing judgments appropriately based on pain experiences, it is not surprising that even representationalists themselves, as mentioned, don't challenge these facts. What is surprising, however, is that they seem not too alarmed by this. Above I continued using the example about feeling pain in my finger. I said: 'Since I am attributing *pain* to my finger, *prima facie* this quality is pain. But if representationalism is true, this quality *is* in fact some physical (extra-mental) condition of my finger.' Call this physical condition Disturbance (*D*). Representationalism implies that the pain attributed to bodily location $L = D$. This implication is empirically falsified. They might still think that the pain experiences themselves (non-conceptually and

²³ It is perhaps worth noting that this point is almost explicitly stated in a note appended to the definition of pain officially recognized by the International Association for the Study of Pain (IASP) — the largest and most influential organization of pain researchers in the world. For references and further discussion, see my (2009, forthcoming-a and forthcoming-b).

²⁴ Note, however, that this is not to deny that LOCATION or FOCUS is true regarding pain experiences. Pain experiences are still transparent in that they are consistent with the Datum. I will come back to this issue below in §6.2.

somatosensorily) represent D in the way in which visual experiences of red (non-conceptually and visually) represent, say, the relevant set of surface spectral reflectances (SSR_R), and grant that our routine pain-attributing judgments are indeed introspective, and as such, don't get to be falsified in virtue of the fact that the perceptual content (that there is D in L) introspectively attributed to the experience happens to be false. If so, they would have to claim that the standard forms of pain-attributing sentences (expressing the relevant judgments) are misleading: they mislead because they incorrectly suggest that a pain/disturbance attribution *is* being made to L . None such is made by these sentences/judgments — these judgments being always introspective, not perceptual. The concept of pain, the representationalists might continue, is not the concept of disturbance (not even extensionally): the former applies to experiences, the latter to what these experiences represent.

But now we have a mystery: we are missing an experientially acquired concept with *de re* labeling uses. Following the parallelism with vision, in the case of seeing red we have the experientially acquired perceptual concept RED that applies to SSR_R , which, according to representationalists, *enables* us to introspect our visual experiences of red by making it possible to judge what it is that we visually experience. Given representationalism and the strong transparency thesis it implies, there is no other way: the relevant range of experientially acquired perceptual concepts with *de re* labeling uses is required for introspection (not, of course, for having the experiences themselves). But then it is a mystery how people routinely can experience pain and introspectively judge/report they are having pain without a corresponding perceptual concept with *de re* labeling uses that apply to disturbances. For clearly they do, not only some of them, not occasionally or rarely. But *all*, who are capable of feeling pain and of judging or reporting they feel pain, as a matter of empirical fact, *do* — and *frequently*. And they do that without, it seems, having a perceptual concept with *de re* labelling uses that applies to D , that according representationalism is represented by the pain experience. So again (P1) is just false.

Is it plausible for a representationalist to respond in the following way? Look, one might say, we do in fact have a concept (or, a range of concepts) with *de re* labelling uses that we apply to D on the basis of pain experiences. This is the concept of disturbance (DISTURBANCE or DAMAGE). I cut my finger while chopping onions, I feel pain there,

and on the basis of this experience I judge that this is tissue damage, referring to the cut. This is plausibly a *de re* judgment made on the basis of my pain experience attributing *D* to my finger with my concept DAMAGE.

But this response won't do. For one thing, the concept DAMAGE is not an sensorially/experientially acquired concept like RED or SWEET — or PAIN for that matter. Its acquisition and experience-based application requires more than the somatosensory or nociceptive modality. It requires more descriptive information than is typically supplied by the somatosensory or nociceptive systems (not only that, it probably, requires historical, biological, and perhaps even cultural or social information). It is clearly not a sensory concept. But, secondly and more importantly, the difficulty that the representationalist faces cannot be resolved by pointing to the *availability* of *some* concept that can be applied directly to the object of pain experiences that, according to representationalism, happens to correctly describe the content of these experiences (if representationalism is true). To resolve the difficulty, the claim must be that it is this concept that is *actually* being deployed when people judge that they are having pain in body parts. For if these pain-attributing judgments are introspective as granted, then according to representationalism these judgments do actually involve the concept that *in fact* describes the content (i.e., *D*) of pain experiences. But this claim is just empirically false. Tons of people correctly make pain-attributing judgments without even thinking about tissue damage or bodily disturbance — indeed without knowing what tissue damage is, or without even having the concept DAMAGE or the like (indeed this seems to be the case with itches and the like, see the next section). Furthermore, it would still remain a mystery why we don't make *de re* perceptual judgments deploying DAMAGE (or the like) as often as we should when we feel pain in body parts, especially in cases where we have either limited epistemic access or don't have any access to the relevant body parts *except somatosensorily or only nociceptively*. Clearly, then, as a matter of empirical fact, we don't need any concept of damage or disturbance and the like to feel pain or to judge/report that we feel pain — and often we don't in fact deploy such concepts when we do judge and report that we feel pain in a body part.²⁵

²⁵ A representationalist might still be tempted by the following (cf. Tye 2006b). We have in fact two concepts of pain, PAIN_E and PAIN_O. The former applies to pain experiences. The latter applies to bodily parts where

5 Other Counterexamples

There are plenty of other experiences that falsify (P1) and (P2). Take itches, for instance. Here is a very typical scenario. I am in front of my computer trying to finish writing this paper, suddenly a particular spot on my back starts to itch for no apparent reason at all. So I am feeling an itch in the middle of my back. Feeling an itch is an experience. I know that I am now feeling an itch. I have introspective knowledge of this experience. How? According to representationalism, I must have the concepts that would express what my experience represents so that I can judge that I am experiencing *that* as ... what? Here 'that' refers to the location in my back (where I feel the itch). This is an extra-mental particular, a bodily region in space-time, and my experience must be attributing an objective quality to that region if representationalism is true. For according to representationalists, this quality cannot be other than a physical condition of that part of my body, and I must have a concept that applies to that location in virtue of its being in the condition represented by my experience. What is this condition and what is the concept I apply to it? The obvious candidate is the concept of an itch, i.e., ITCH. What objective property does it attribute? I have no idea. In fact, I have no idea what objective property I am *perceiving* — let alone representing. When pressed like this, representationalists use the same maneuver likening itches to the perception of secondary qualities: when experiencing red, just as I may not know what complex physical property red is that I am perceiving except that it is some physical condition of surfaces distinct from this or that colour, etc., I may not know what physical condition ITCH attributes except that it is some condition different than tickles, somewhat similar to pains, etc. — when I judge I feel an itch in my back. It may be that science will discover the nature of itches occurring in

we feel pain. Sometimes I deploy PAIN_O when I judge that I feel pain in my finger: my judgment is correct iff my finger is represented by my PAIN_E to be damaged. That is, my finger has pain in it when and only when I experientially represent it to be damaged. So the concept PAIN_O attributes an inverse intentional (therefore, mental, subjective) property to my finger when I judge I feel pain in my finger (also cf. Bain 2007). This is an interesting suggestion. In fact, the positive account I will give later will in some ways be structurally similar to this suggestion (see §6.2 below). But it doesn't save representationalism. Neither concept attributes a property represented by pain experiences. Both concepts, to the extent to which they have *de re* applications appropriately based on pain experiences, track something subjective, experiential, and to that extent, introspectable. This violates S-TRANSPARENCY. The features to which these concepts apply, or indeed the features/qualities they may express, are introspectable features of experiences over and above those implicated in their representational content.

bodily locations, just as science has told us that red is (let us assume) a set of surface reflectances of such and such kind.

Unlike in the case of pain, however, where we have years of scientific research, we don't know much about what physical conditions might be correlated (*sic.*) with itch experiences. At any rate, it is empirically evident that attribution of itches to body parts do not get falsified with the presence or absence of any physical conditions of those parts. When we introspect itch experiences, we are presented with some features for which there are no concepts with *de re* labeling uses applying to extra-mental conditions. Hence, itch experiences are not strongly transparent.

Take orgasms. Correctly judging that one is having an orgasm is judging that one is having an experience of a certain sort. Thus, it is an introspective judgment. Evidently, the concept ORGASM does not, *de facto*, have any *de re* labeling uses such that reporting an orgasm is correct or incorrect according to whatever extra-mental conditions obtain in the relevant parts of one's body. More perspicuously, the essential (introspectable) phenomenal features that make an experience an orgasm rather than some other kind of experience are features for which there are no concepts with *de re* labeling uses. Hence, introspection of orgasms *qua* orgasm does not rely on concepts that attribute extra-mental features or conditions that orgasm experiences may represent relevant bodily parts as having (if they represent any such things). Thus, orgasms (i.e., experiences of a certain sort) are not strongly transparent, therefore not representational.²⁶

What about affective/hedonic experiences that fill our daily lives? Indeed, take the *painfulness* of pains, or the impressive explosive *pleasantness* of orgasms — although any other pleasant or unpleasant sensory experience could function as a counterexample (such as the pleasantness or unpleasantness of many taste, smell, auditory, tactile experiences we have on many different occasions). Phenomenal qualities such as pleasantness or unpleasantness, insofar as they are phenomenologically distinct and salient, *primarily* qualify experiences and not the extra-mental objects or particulars that these experiences may represent. If we sometimes attribute affective qualities to the objects of our

²⁶ Tye (1995, 1996b), responding to Block's challenge (1996), argues that orgasm experiences are representational and thus strongly transparent. The arguments he marshals for this claim, however, establish *at best* that these experiences are intentional, not strongly representational. Please note that none of the arguments I present here against representationalism is meant to be an argument against intentionalism *per se*.

experiences, as we clearly do, this depends on our prior understanding of whether the experiences themselves have (generally) the affective/hedonic qualities in the first place.²⁷ But even when we attribute affective qualities to the extra-mental particulars through our experiences, the concepts such as PLEASANT and UNPLEASANT (just like the concepts PAIN, ORGASM, ITCH, etc.) do not have *de re* labeling uses: our *de re* judgments about these particulars do not get to be verified or falsified on the basis of whether these particulars have certain extra-mental features or in certain physical conditions. Thus affective experiences *qua* affective are not strongly transparent, thus not representational. Pleasantness of orgasms or painfulness of pains do not represent extra-mental conditions of body parts.²⁸

The same line of reasoning will similarly give us the result that emotions and moods, or at least those experiential aspects of these states that are phenomenologically real and distinctive, are not strongly transparent, thus not representational. Just think, for instance, whether the experiential *feel* of anxiety, grief, sadness, euphoria, etc., are strongly transparent in the sense required by representationalism.²⁹ I trust that at this point there is no need for me to explicitly spell this reasoning out — the pattern should be clear by now given the empirical facts about how our introspective judgments about these psychological states work.

²⁷ For a straightforward argument to this effect, see Aydede & Fulkerson (2014). Representationalists sometimes are tempted to provide an explanation appealing to the affect of experiences (their unpleasantness or pleasantness) for why we do not have the relevant range of *de re* labeling concepts for intransitive bodily sensations. It is not clear whether such an explanation would be correct or entirely correct: note that there are affectively neutral or nearly neutral versions of these sensations — this is in fact the more pervasive norm — where the phenomenon still persists. But, more importantly, even if such an explanation were correct, it would not save the representationalists. For it would be an *admission* that we have violations of S-TRANSPARENCY even if there may be good naturalistic reasons for the practice. Finally, insofar as the affective aspect of experiences is a phenomenological matter, the ‘explanation’ offered would itself constitute a refutation of representationalism if the affect primarily qualifies the experiences themselves. For the explanation would amount to admitting that there are intrinsic phenomenal/affective features of experiences that are introspectively available. Note that at this point the representationalist cannot argue that the affect itself is strongly transparent. See Aydede (2009) and Aydede & Fulkerson (2014) for further details.

²⁸ Some representationalists such as Bain (2013) and O’Sullivan & Schroer (2012) think that the properties represented by the hedonic valence of experiences (in particular, the painfulness of pains) represent evaluative properties such as the goodness or badness of extra-mental objects or conditions represented by affective experiences. I will put these proposals aside since their naturalistic credentials are moot. I have criticized such views elsewhere — see Aydede (2006) and Aydede & Fulkerson (2014, forthcoming)

²⁹ In fact, as mentioned before, it is not clear whether these experiences are even transparent at all, strongly or otherwise. See Kind (2013) for a persuasive argument that they are not.

Note that above cases are even less amenable to the kind of response given by a representationalist in response to my parallel claims above in the case of pain experiences. In the case of pain, at least we have plausible candidates (tissue damage, disturbance, etc.) popularly associated with pains. With these other examples discussed in this section, we don't even have obvious candidates (popular or otherwise) that would serve as the represented content of these experiences so that the claim that their concepts can be put in *de re* labelling uses has much plausibility. At any rate, if pressed, my rebuttal would parallel the one I made above in discussing pain.

I conclude that there are kinds of experiences that disconfirm P1 and P2. Therefore S-TRANSPARENCY, thus representationalism, is empirically false. The experience of pain is simply the most salient and philosophically the more widely discussed experience among many such 'intransitive' experiences that are not strongly transparent.³⁰

In the next section, I will propose a framework about how introspection of perceptual experiences works that will make sense of the Transparency Datum in light of the rejection of strong transparency.

6 Introspecting the Phenomenal Qualities of Experiences

The strong transparency thesis (S-TRANSPARENCY) claims that experiences with phenomenal character have no introspectable features over and above those implicated in their representational contents. I have just argued that there are plenty of experiences with a robust and distinctive phenomenology that make this claim false. So some experiences have introspectable features over and above those implicated in their representational content. More perspicuously, some experiences have introspectable phenomenal features that cannot be attributed, as representational content, to the aspects of the extra-mental world. Still, I also think that this is not in conflict with the Transparency Datum which is the conjunction of two observations supported by introspection: LOCATION and FOCUS. LOCATION roughly says that the qualities that we become aware of in having an experience are qualities that are presented to us as belonging to extra-mental particulars. FOCUS says

³⁰ The term 'intransitive' is from Armstrong (1961, 1968) who uses the term to mark only the subcategory of bodily sensations (pains, itches, tickles, etc.) that he notes are problematic for perceptual theories for reasons similar to ones raised here.

roughly that it seems impossible or at least is very difficult to attend to phenomenal qualities *of* an experience without attending to the extra-mental particulars that one's experience presents one as having these qualities.

The rejection of S-TRANSPARENCY and the acceptance of the Datum create a *prima facie* tension that needs addressing. If pain experiences falsify S-TRANSPARENCY, how can they be consistent with the Datum? Without some sense of how this tension can be assuaged, one might just think that if the Datum is granted, S-TRANSPARENCY is worth arguing for, especially if it otherwise saves us a lot of philosophical headache on various grounds.³¹ Alas, completely removing this tension requires developing a substantive account of introspection that explains one's epistemic access to one's experiences. I have developed such an account of phenomenal introspection elsewhere (Aydede & Güzeldere 2005) — a detailed account that is both intentionalist and naturalist. Here I want to extend that framework (with some corrections added) to give an explanation of how the denial of S-TRANSPARENCY is compatible with the Datum.

6.1 General Framework

Let us focus on sensory and perceptual experiences (including bodily sensations) and leave aside moods as well as affective and emotional experiences. Perceptual experiences have an intentional structure — they represent. Experiences, in other words, have both a referential and a predicative structure. In perceptual experience, we are presented with particulars having properties and relations. Thus our perceptual systems have ways of picking out or referring to token objects, events, states, space-time points or regions — particulars, in short, very broadly understood as tokens in space-time. Researchers working on perception have postulated various mechanisms in which these ways of referring can be

³¹ Indeed, there are philosophers who think that there are very powerful independent theoretical and methodological reasons to think that perceptualism or representationalism is true. So they tend to think that if perceptualism/representationalism is true, then the ordinary as well as scientific conception of pain and other intransitive bodily sensations is just incoherent (see, for instance, Hill 2006, 2009). This is puzzling, on independent grounds: if the concept of pain with which the scientists have been operating were incoherent, we should be seeing the troubling signs of this in both the basic scientific research at large and clinical practice. But as far as I can tell, none of the kind exist. In fact, both have been exponentially flourishing after the scientific revolution the pain science witnessed in the 1960's, which prompted the IASP to adopt a definition of pain that in fact embraced the folk conception of pain which rejects identifying or even robustly correlating pain with tissue damage or the like — for a critical discussion of the IASP definition of pain, see my (forthcoming-b).

implemented — they usually take the form of topographic map-like representational vehicles that map onto one’s sensory fields, e.g., one’s visual, auditory, somatosensory fields, and master maps that coordinate these, etc. But as is well known, reference is not sufficient for representation with correctness conditions. A predicative structure that attributes features, qualities, properties, and relations to particulars is also necessary. Even experiences generated by what is traditionally thought to be informationally poorer modalities such as olfaction and gustation attribute properties (e.g., bitterness, sweetness, pungency, etc.) to particulars such as the odour molecules one is sniffing now or the substances in contact with one’s tongue.³² Vision attributes properties to particulars in one’s visual field, various bodily senses attribute qualities to particular points, regions, or volumes in one’s body (one’s somatosensory field). Sounds are similarly attributed to particulars in one’s auditory field. These are fairly general remarks about the intentional structure and organization of perceptual experiences. None of this should be very controversial.

It is plausible to think that if the phenomenality of perceptual experiences is philosophically mysterious and puzzling, this is mostly because of the predicative (not referential) structure of such experiences. Roughly, the idea is that the puzzle of phenomenal (experiential) consciousness can almost exclusively be traced to the *attributive* (predicative) function of experiences. In particular, it is, for the most part, the perceptual attribution of so-called secondary qualities to particulars presented in one’s experience that seems to generate the puzzle. I take secondary qualities to be complex physical properties whose detection is modality specific.³³ Nevertheless, the phenomenology of the perceptual attribution of such properties to particulars does not give one epistemic access to the physical nature and complexities of these properties. They are not presented to us in experience as having much complexity at all. In fact, the positive, specific and substantive

³² Batty (2010) argues that the content of olfactory experience is not referential but quantificational and the only reference to particulars are via general indexicals such as *here, now*. I am not sure I agree with her analysis, but even if her argument is sound, this will not pose any problem for my analysis below as my emphasis will be on the predicative structure of experiences and their introspective expression. Furthermore, if she were right, we would have more support for rejecting S-TRANSPARENCY. See Mole (2010) for a criticism of Batty.

³³ This claim is not, strictly speaking, necessary for the purposes of this paper. But in what follows, I will assume a primary quality (physicalist) view of secondary qualities.

phenomenology of sensory experiences suggests that these properties *are* presented more or less *as simples*. I will take this phenomenology at its face value and assume that each sensory modality has its own distinctive predicative system that determines the phenomenal quality space defined for that modality and its proper sensibles.³⁴ So the formal structure of sensory experiences takes the form of:

#*x* is *F*#

where #*x*# refers to all particulars in one's sensory field and #*F*# attributes a proper sensible to those particulars detectable in the sensory modality responsible for the experience.³⁵ I will take the sensory predicate #*F*# to be part of the representational vehicle implementing the experience. As such, #*F*# itself is (or determines) a fully determinate phenomenal quality located in the quality space specific to the modality in question. Despite the displayed format, the intentional structure of the experience is not conceptual (to give the flavour: the representational vehicles implementing the perceptual experiences are more like maps than sentences). To fully specify the non-conceptual content, one needs to specify which sensible property is attributed to each *x* located in one's sensory field.³⁶

³⁴ For the notion of the quality space for sensory modalities and the science behind it, see Clark (1996, 2000). Although I assume that there is a quality space for each sensory modality, this assumption is not necessary for my purposes here. See Aydede & Güzeldere (2005) and Kulvicki (2004, 2005) for elaboration of how certain physically complex properties could be phenomenologically presented as if they were simple — or at least how phenomenology could be silent about the physical complexity of secondary qualities. The point goes back to Smart (1959) and Armstrong (1968) — for a similar recent development of the idea, see Fazekas (2012).

³⁵ I will use '#...#' to refer to the vehicle of sensory representation. So, for instance, #*F*# (the sensory predicate) attributes the sensible property *F* to an extra-mental particular, *x*, picked out by #*x*#.

³⁶ As mentioned before, I will leave aside the elaboration of how '*F*' can range over properties that are high-level and multi-modal such as shape, size, motion, or even some physical kind properties such as being a pine tree. Here I just want to concentrate on low-level proper sensibles as I think most philosophical puzzles about perceptual phenomenology stem from *their* peculiar phenomenology.

Also, I leave the exact nature of particulars deliberately vague as they can be token objects, events, as well as space-time points or regions — a more careful discussion would involve distinguishing space-time points as providing the sort of 'Kantian' scaffolding for sensory representation and the particular items (objects, events, etc.) occupying this space-time. The point of experiences having a singular referential structure is to point out the obvious: the world our experiences disclose is a world of particulars in space-time instantiating properties. So the representational resources of such experiences are capable of specifying the spatiotemporal distribution of sensible properties that it attributes to particulars in this space-time. This requires that experiences have *referential* as well as attributive functions. See Peacocke (1992), Clark (2000), Burge (2010).

Perceptual experience supplies us with information about the extra-mental particulars and *their* features. When we introspect such experiences, this is what we find, just as LOCATION says: we find features experienced as instantiated by extra-mental particulars, and our introspective knowledge is articulated (partly) by concepts that apply in the first place to these particulars. One way to read S-TRANSPARENCY is as claiming that introspective knowledge of experiences is exhausted by whatever can be expressed by the self-attribution of a complex predicate of the canonical form:

(FORM-1) EXP(x is F)

where ' x ' and ' F ' are meant to indicate referential and predicative positions respectively in the self-attribution of representational content as experienced, where both ' x ' and ' F ' range over extra-mental particulars and properties respectively.³⁷

A way of seeing the motivation behind the internalist denial of S-TRANSPARENCY is to wonder whether there is more to introspective knowledge than can be expressed by FORM-1. Suppose you ask me how this tomato looks to me in regard to colour.³⁸ I report correctly by saying that it looks red to me. You say, 'Yes, yes, I know, but how does its

In a very sketchy form, however, my view is that perception (at least visual perception), in presenting the particulars it does, uses at a minimum, in addition to attributives/predicates, a referential device based on a spatial grid-like (functional) topography built into the vehicle of perception. The idea can be illustrated in analogy (and only in *analogy*) to distinct points on the display of a digital camera picking out different spatial locations in front of the camera. Here there is a pre-established isomorphism (under suitable conditions) between the topography of the display and the space in front of it (sustained by optical laws and geometry) and this isomorphic correspondence can ground *reference*. Property *attribution* (predication) would then consist of the systematic *causal* correspondence between two sets of property instantiations at these points/locations (correspondence between properties instantiated on the display and the properties instantiated at the locations in front of the display) — an informational psychosemantics could then ground the semantics of perceptual predicates. As mentioned, there may in fact be more than one referential scheme — for instance, schemes corresponding to representing space-time points/regions and representing particular objects/events occupying these points/regions (cf. Pylyshyn's visual indexes or FINST's, 2007). But these are mostly empirical matters whose discussion needs some other occasion.

³⁷ More colloquially: 'I am experiencing x as F ', which can be paraphrased as ' x looks/appears F to me'. I will sometimes use this form. Reference and predication can take demonstrative forms as in 'I am experiencing *this* as *that*' where 'this' picks out a particular and 'that' attributes a quality specified demonstratively. This is usually how we achieve communicating richer and more determinate content than can be expressed by our standing non-demonstrative concepts. I will leave this aside. Also, the predicate 'EXP' can, of course, have more specific forms such as 'SEE', 'HEAR', 'FEEL', etc.

³⁸ The following dialogs are inspired by Lycan (1996: 124; similar insightful dialogs appear in his other writings).

looking red feel to you intrinsically?’ (or, ‘How is its redness experienced by you?’). If there is a legitimate non-trivial answer to this latter question, then it cannot be given in the form allowed by (FORM-1). We can generalize:³⁹

You ask: How do you experience this? (How does this look to you?)

I answer: I experience it as *F*. (It looks/appears *F* to me.)

You: Can you tell me the *way* you experience this as *F*? (The *way* it looks *F* to you?)

I: Uhm... !?

This last question is the question of asking *what it is like* for me to experience this as *F*.⁴⁰ Here the form of a proper answer would be something like:

(FORM-2) It is like *Q* for me to experience *F* (or, the *F*ness of this),

where ‘*Q*’ is meant to express a concept that I use in introspection specifying the *way* I experience *F*ness of a particular that is demonstrated. Note that the question presses on the way *F*ness is experienced by me, not on how *this* being *F* is experienced by me — it is a question about how the *way* the *predicative* structure of my experience presents the sensible property it does. Call concepts meant to specify, classify, or categorize (or, in this sense, meant to attribute a quality to) the *ways* that the sensible properties of extra-mental particulars are experientially presented, *phenomenal concepts*, and the introspective knowledge meant to be expressed by their use in judgment of FORM-2, *phenomenal knowledge*. Clearly, if phenomenal concepts exist, we do not seem to have natural language terms to express them (except perhaps in the case of intransitive bodily sensations — see below).

³⁹ Here again ‘this’ refers to an extra-mental particular, and *F* is a sensible property of this particular.

⁴⁰ Note the difficulties confronted by all representationalists responding to Jackson’s Knowledge Argument that requires a robust sense of introspective knowledge of what-it-is-like to experience *F*. FORM-1 cannot capture this sense — witness Dretske’s struggle in his (1995, Chapter 3).

Phenomenal knowledge is introspective knowledge of features of experiences that are thus over and above those implicated in their representational content. *Understood this way*, representationalists, when they are careful, ought to reject that we have phenomenal knowledge of our experiences. By contrast, phenomenal internalists defend the view that we can have (and many times, do in fact have) phenomenal knowledge. Remember that phenomenal internalists need not reject intentionalism about phenomenal character. On the contrary, since phenomenal knowledge is introspective knowledge of *the way sensible properties are presented to one in experience*, there is a natural understanding of intentionalism according to which phenomenal knowledge is introspective knowledge about intentional facts.

Do we have phenomenal knowledge at all? Well, if S-TRANSPARENCY is false, then we do. This is one of the main conclusions of this paper: If we know we have pains, itches, orgasms, unpleasant experiences, etc., then we have phenomenal knowledge as specified. I will come back to this shortly. For the moment, let us continue to discuss phenomenal knowledge we can have about regular perceptual experiences.

So, how is the acquisition of phenomenal knowledge possible given the Transparency Datum — compatibly with LOCATION and FOCUS? A phenomenal internalist needs to be able to make sense of this. To begin with, note that LOCATION says something about the intentional structure of experiences: experiences have both referential and attributive functions — they present to us both extra-mental particulars and *their* properties. There is no problem when we report this using the introspective FORM-1. Phenomenal knowledge, however, through whatever introspective mechanisms it may be delivered, is only attributive (or better: second-order quantificational): it says *of some way in which Fness is presented* to me in my experience that it is *Q*. FORM-2 does not have a singular referential position; it does not pick out a particular apart from the one whose sensible property it experientially registers. This is why it has been so natural to use the ‘what-it-is-like’ construction in the expression of phenomenal knowledge.

So, suppose I am experientially aware of a particular as *F*. When I make an introspective judgment of FORM-2, the only particular I am aware of, then, is the extra-mental particular that my experience presents to me whose *Fness* is experienced *in a certain way* — as *Q* or *Q*-ly. Put differently, my experience presents a certain extra-mental

particular to me as *F*, and then, my phenomenal knowledge consists of my applying a phenomenal concept to an ‘object’ conceived by me only *as the way Fness is experientially/perceptually presented to me now*. I am in no way aware of this ‘object’ in the sense in which I am aware of extra-mental particulars in sensory/perceptual experiences. In short, I do not sense, perceive, or in any other way experience, this ‘object’, i.e., the way *Fness* is presented to me in my experience. Introspective mechanisms do not have the vehicular resources to make demonstrative singular reference to experiences. Nevertheless, their predicative resources are distinctive and give us direct grasp of phenomenal properties as certain ways in which the sensible properties of extra-mental particulars are presented to us in experience.⁴¹ The deliverances of introspection is immediately conceptual. Introspective mechanisms are not sensory mechanisms. Introspection does not have its own proprietary ‘introspective field’ where mental particulars are located and displayed. So whatever sense in which we can attend to or focus on extra-mental particulars in sensory experience, this is *not* the sense in which we can attend to or focus on the phenomenal qualities and their instantiations. Nevertheless, if, in whatever sense, we want to *introspectively* ‘attend to’ the way in which the sensible properties of this or that particular are presented to us, we cannot help but *perceptually* attend to these particulars whose sensible properties are being registered in our experience in certain ways. In other words, we can *conceptually* attend to phenomenal properties only by way of *perceptually* attending to the extra-mental particulars whose sensible properties are sensorially registered in our experiences in certain ways — these ways constituting (or, determining) the phenomenal qualities of our experiences. This would explain why

⁴¹ I would not mind putting this point by saying that although we can have direct introspective knowledge of the instantiation of phenomenal properties, this is accomplished without the mental equivalent of a demonstrative singular reference to the state or event that does instantiate those properties. This is probably what direct acquaintance with phenomenal qualities comes to. But, unlike Bertrand Russell, I do not think such direct acquaintance would enable one to demonstratively pick out one’s experience whose qualities one is said to be acquainted with. Of course, internalists sometimes express their knowledge by locutions like ‘*this* is what it is like to experience *F*’ in English. But this is fine. There is not much else that could be done with the resources of natural languages — although here ‘this’ could naturally be interpreted as referring merely to the particular instantiation of a phenomenal quality to pick out its kind. There is, of course, a more natural and widely used locution to express one’s phenomenal knowledge: ‘This looks *F* to me’. This locution and the like, when used in the phenomenal sense (see Chisholm 1957, Jackson 1977), does a very good job of expressing one’s phenomenal knowledge. To mark this sense and avoid some issues with the semantics of ‘*F*’, we can use a hyphen: ‘this looks-*F* to me’ where ‘this’ refers to an extra-mental particular.

LOCATION and FOCUS strike us as phenomenological truths while at the same time allowing us to deny S-TRANSPARENCY.⁴²

Note that we do not have any difficulty of sensorially imagining secondary qualities with which we have been acquainted. If I ask you to vividly imagine a clear and intensely blue sky on a sunny spring day and then ask you to concentrate on its colour, it makes little sense to ask you *where* the colour quality you are imagining is represented as being instantiated — despite the fact that there is usually a very strong sensation-like colour phenomenology in your experience that you are introspectively aware of. So we are perfectly comfortable with the idea of having introspective knowledge of the *way* in which a proper sensible is (re)presented in our experience without this imagined quality being attributed to any particular space-time region, and for that matter, without this *way* being attributed to any particular, mental or otherwise.

6.2 *Framework Applied to Pain*

Now let us go back to pain experiences. We have observed that sincere judgments locating pains in body parts come out true whether or not there is anything physically different (disturbance) in those parts. We have concluded that our pain-attributing judgments are

⁴² I have generally tried to avoid the use of ‘qualia’ in this paper. But if we wanted to associate this story with a story told in qualia terms, here is how it would go. A quale type can be identified with *the* or *a way* of experientially registering a sensible property for which a quality space can in principle be specified. If *Q1* is a way of registering red_{16} for a certain subject, it is metaphysically (or perhaps, even empirically) possible for another subject to register red_{16} as *Q2* ($\neq Q1$) — according to some matrix specified by some quality space defined for these *Qs*. So this view would allow for shifted or inverted qualia. On this view, qualia play a role similar to the role Fregean contents play (see, for instance, Chalmers 2004, Thompson 2009). But qualia are not contents in my view. A particular quale type is a *sensory* predicate type deployed in sensory experiences whose concept is a phenomenal concept (another predicate — a *conceptual* one) identifying this type. Qualia, of course, normally attribute sensible properties to extra-mental objects, but sometimes not as in intransitive sensations (if we adopt my second story about pain quality in the main text — see below). So there are no Fregean contents in any traditional sense. But it is possible to identify these predicates as ‘modes of presentations’ (MoPs) of the properties they attribute (when they do). But this is a degenerate sense of MoP, since we might as well just talk of predicates syntactically typed whose content is just Russellian (they express extra-mental properties — when they do — directly, unmediated by senses). Unlike Fregean contents, qualia do not determine which properties they express. This is determined by some informational psychosemantics — in my view, qualia are just representational vehicles with a certain functional/information-theoretic role. For details, see Aydede & Güzeldere (2005). Clearly, the view advocated here has very close affinities with the qualia friendly adverbialist views of perception developed in the 60’s and 70’s, but it does not suffer from the devastating problems those views generally thought to have. For an explicitly adverbialist account of sensory affect, see Aydede (2014). Alter (2006) gives an account of phenomenal manners of representing in experience that is similar to the account given here — although he doesn’t make a reference/predication distinction, he seems to have predication in mind.

introspective rather than perceptual. Nothing of this sort happens, for instance, in case the surface we visually experience as red turns out to be white (say, when seen under tricky conditions). We do not classify our judgment attributing the colour of red to the surface as introspective. This judgment ('that surface is red') remains a *de re* perceptual judgment — it is just false. Hence there is an appearance/reality distinction for colours (just as there is for other secondary qualities indeed). So how do pain-attributing judgments play out according to the framework developed above?

Let us focus on the pain *experiences* first. These experiences are sensory with a somatosensory field where extra-mental particulars (bodily regions) are represented to be located. There is also a quality space whose coordinates determine the sensible properties to be attributed to the bodily locations. Hence, pain experiences, just like other sensory experiences, have a referential and a predicative structure working together. These experiences feed into a conceptual system wherein introspective and perceptual *judgments* are made based on these experiences. The judgments made locating pains in bodily locations track what sensory predicates are deployed, not what sensible properties/conditions are thereby attributed to these locations by these predicates. Thus, whatever property is attributed to bodily locations, our judgments are about the *ways* these properties are experienced, or sensorially registered — not about the properties these ways attribute. Because these ways are *ways* in which certain properties are sensorially attributed to extra-mental particulars (bodily locations), we cannot help but attend to these *ways* except by attending to the locations instantiating these properties sensorially attributed. These ways are (or, determine) the phenomenal qualities of our experiences, whose knowledge is thus knowledge of the *ways* in which certain conditions are sensorially attributed to body parts. Thus when I judge I have a pain in my finger, my judgment is correct in virtue of my undergoing a pain experience attributing a property to my finger in a certain way. My judgment thus correctly reports an experience — it is an introspective judgment.

What is the property that seems to be attributed by my experience? Here unlike a perceptualist or representationalist, we are not theoretically constrained about what these properties ought to be. In fact, we may *legitimately* draw a blank — just as folk and

scientists do. But there are various options. I will mention three but explore only one of them here a little.

On the first option, following representationalists or perceptualists we might say that this property is some sort of physical disturbance. If we say this, our pain-attributing *judgments* would still come out as correct, as desired, but our *experiences* now may *not* be veridical. I may correctly report pain in my thigh when in fact there is nothing physically wrong with it and my pain is a referred pain due to a pinched nerve in my lower spine. My pain *experience* is thus illusory but my pain *judgment* is still correct. I suppose we can learn to live with this result — even though, as I have argued, the lack of relevant *de re* perceptual judgments would make these experiences not perceptual or (strongly) representational. Furthermore, it is not clear whether this proposal does justice to the phenomenology of locating sensations in body parts while resisting to the identification of it with a physical property.

A second option is to say that the ‘properties’ that the sensory predicates seem to attribute to body parts are *mental* (or, mind-dependent) in the following sense: a body part has pain in it just in case it is the intentional target of a sensory predicate predicating a dummy property of that location. The pain experience I have due to a paper cut in my finger makes a successful reference to my finger and involves activation of a predicate (or, a range of predicates) in virtue of which it is true that my finger has pain in it. The mental property (pain) that qualifies my finger is the inverse quasi-intentional property of *being the target of a sensory predicate being used with respect to a referential position* (that picks out my finger). Although the reference succeeds (in this case), the predicate activated doesn’t attribute any real property to my finger (hence ‘dummy’, see fn. 44 below). Furthermore, the inverse quasi-intentional property had by my finger is *not* represented by my experience either, but my judgment ‘I have pain in my finger’ is made true because of it. In other words, as long as sensory reference succeeds, the reference (my finger) is guaranteed to have the property attributed insofar as the system does a predication with respect to that reference. This allows us to usefully distinguish between informative pains, referred pains, and phantom limb pains.

Informative pains are those when the predication actually signals or indicates actual or potential physical disturbance/damage at the location to which reference is successfully

being made. (Correlations between physical disturbances and firing of a predicate have been claimed to be fairly poor — but perhaps when the channel conditions are right, there is information flow after all, even though this may not be enough for genuine representation. See Aydede & Güzeldere 2005 for details).

Referred pains are those in which reference is successfully made to actual body parts with respect to which a predicate is *causally activated* by some disturbance in some other part of the body, but the activation does not *indicate* disturbance in the part of the body to which the reference is actually being made — perhaps because the channel conditions are not quite right.

Phantom limb pains are those where attempted reference to a body part fails but with respect to which a sensory predicate is nevertheless activated.⁴³

If the property ‘attributed’ to body parts is mental in this sense, then there are no representational mistakes anywhere in the intentional system. Our experiences are intentional (because there is either successful or failed reference). But the pain experiences do not make genuine property attributions to bodily locations — so they do not have full veridicality conditions.⁴⁴ And our *judgments* prompted by them are correct in intuitively

⁴³ More needs to be said about phantom limb pains. There are extra-special difficulties with phantom limb pains due to reference to absence of limbs in our judgments/reports. I cannot discuss these here. See Bain (2007) for an insightful discussion of the problems in understanding the locations of pain in phantom limbs.

⁴⁴ Consider an apprentice among alchemists in the pre-modern world pointing to the vapor coming out of boiling water. He utters, ‘this is phlogiston’. Given my Russellianism and the fact that there is no property of being phlogiston, the apprentice is not making a genuine property attribution — although his reference is successful. But although his utterance is not strictly speaking true or false, there are nevertheless appropriateness or suitability conditions to his utterance that are not satisfied in this particular case. And that is what would be pointed out to him when his tutors point out his ‘mistake’ — this description is of course from the perspective of a semanticist. Pain experiences, on this option, are *like* this utterance. We might say that they do not make *genuine property* attributions, or we might even say that they do not *genuinely make* property attributions. Either way, they are not fully representational (or perhaps: they are attempted but ‘failed representations’). And this view has been the dominant view pretty much throughout the history of philosophy — witness Reid in the above quotation. Or here is McGinn (among many others): ‘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: 8). For many other references to such traditional views, see Bain (2003: 502).

Another option to characterize the ‘property’ attributed by pain experiences might be provided by Pautz’ view (2010). If I understand him correctly, Pautz takes this property to be a primitive one that lives only in the intentional content of these experiences somehow projected to bodily parts in a way that does not generate accuracy conditions. If this is meant to be consistent with the kind of non-representationalism this second option explores, I am sympathetic (barring my worries about how to naturalize such a content) — however, I am not confident that I got Pautz’ view right.

the right way. These judgments usually correctly locate the mental properties ‘located’ in those bodily regions, when the relevant mental properties are understood in the above way. But these mental properties are not represented in the experiences; as said, pain experiences, although intentional (*de re* reference), are not fully representational (no genuine predication generating accuracy conditions). This is, of course, not to deny that there *is* a quality (pain) instantiated in my finger — it is just to deny that this quality is represented in my experience — nevertheless, this quality is what my judgment attributes to my finger. Folk and the pain scientists (including clinicians) do routinely attribute sensations to bodily locations after all — come to think of it: *sensations* in extra-mental particulars. This story about how to account for the ‘properties’ that seem to be sensorially attributed to body parts in pain experiences explains how and why.⁴⁵

A third option may be provided by various forms of projectivism, according to which the properties attributed to body parts by pain experiences are in fact qualities of the experiences themselves, or qualities somehow associated with these experiential qualities. On these views, we somehow systematically project these qualities to body parts that actually belong to, or associated with, the pain experiences.

There may be other options.⁴⁶ Although I am officially neutral on which options to take, I am leaning toward exploring the second option more. But I will leave the discussion of these options, their strengths and weaknesses, to some other time. What is important to keep in mind is that there *are* a lot of ways of making sense of how S-TRANSPARENCY can be denied consistently with the Transparency Datum (and a robust sense of realism about phenomenal properties or qualia), both in standard perceptual experiences at large and in intransitive bodily sensations such as pain, without compromising a naturalistic framework, such as above, within which various engineering solutions can be explored.

⁴⁵ Note that this is *not* a form of projectivism — there are no representational mistakes anywhere in the system. One robust mark of projectivism is that it makes experiences under consideration and our judgments based on them massively illusory or somehow mistaken.

⁴⁶ For instance, the properties experientially attributed to body parts may have the form of Shoemaker’s ‘appearance properties’ — see his (1994 & 2000). For a useful discussion and comparison, see Block (2006).

7 Summary and Conclusion

I have argued that empirical facts about pain experiences and the way we conceptually respond to such experiences falsify S-TRANSPARENCY. Thus, representationalism that implies S-TRANSPARENCY is false. Nevertheless, I have shown how we can respect facts about experiential transparency given a naturalistic account of introspection that allows for unmediated epistemic access to phenomenal properties of experiences.

Pains and other intransitive bodily sensations are tough phenomena for a theorist of any stripe, since the way *everybody* thinks about them (including the representationalists themselves) is as what they are, as *sensations*. We *both* locate them in body parts in the extra-mental world, *and* at the same time, treat them as essentially mind-dependent (as sensations) — not as a matter of theory, philosophical or otherwise, and not in laboratories or clinical settings either, but as a matter of routine pretty much *anywhere*. Because of this, they are particularly tough for representationalists. If representationalism is true, they ought not to exist. Call for their elimination is occasionally heard on the basis that the concept of such sensations is incoherent.⁴⁷ But it is not at all clear how seriously we can take such calls since even the defenders of the incoherence view do not think that we can in fact get rid of our concept of pain, and they acknowledge that scientific theorizing and clinical practice show no sign of trouble due to our concept of pain. So, such calls are mostly the result of ideological posturing. I suggest that we start taking pains and other similar experiences seriously for what they show in general about phenomenal consciousness itself.⁴⁸

References

Alter, T. (2007). ‘Does Representationalism Undermine the Knowledge Argument?’ In *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, by Torin Alter and Sven Walter, Oxford University Press, 2007, 65–75.

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

⁴⁷ See, for instance, Dennett (1978), Hardcastle (1999), Hill (2006, 2009).

⁴⁸ Many thanks to John Kulvicki, Tyler Burge, Matthew Fulkerson, Brendan O’Sullivan, David Bain, and Michael Brady for useful feedback on earlier versions of this paper. Thanks also to my colleagues, Evan Thompson and Adam Morton, for their insightful comments and questions during a departmental “work-in-progress” session in which parts of this paper were discussed.

- Armstrong, D. M. (1968). *A Materialist Theory of the Mind*. New York, Humanities Press.
- Aydede, M. (2002). “Is Introspection Inferential?” In *Privileged Access: Philosophical Accounts of Self-Knowledge*, Brie Gertler (Editor), Ashgate Publishing (*Epistemology and Mind Series*).
- Aydede, M. (2006). “The Main Difficulty with Pain.” In *Pain: New Essays on Its Nature and the Methodology of Its Study*. M. Aydede, Ed. Cambridge, Massachusetts: MIT Press.
- Aydede, M. (2009). “Is Feeling Pain the Perception of Something?”, *Journal of Philosophy*, 106(10): 531–67.
- Aydede, M. (2014). “How to Unify Theories of Sensory Pleasure: An Adverbialist Proposal,” *Review of Philosophy and Psychology*, 5(1): 119–33.
- Aydede, M. (forthcoming-a). “Defending the IASP Definition of Pain,” *The Monist*.
- Aydede, M. (forthcoming-b). “Pain: Perception or Introspection?” In *The Routledge Handbook of Philosophy of Pain*, edited by Jennifer Corns, Routledge.
- Aydede, M. and Fulkerson, M. (2014). “Affect: Representationalists’ Headache,” *Philosophical Studies*, 170(2): 175–198.
- Aydede, M. and Fulkerson, M. (forthcoming). “Reasons and Theories of Sensory Affect.” In *The Nature of Pain*, edited by David Bain, Michael Brady and Jennifer Corns, Oxford University Press.
- Aydede, M. and Güzeldere, G. (2005). “Cognitive Architecture, Concepts, and Introspection: An Information-Theoretic Solution to the Problem of Phenomenal Consciousness,” *Nouûs*, 39(2): 197–255.
- Bain, D. (2003). ‘Intentionalism and Pain,’ *Philosophical Quarterly*, 53(213): 502–522.
- Bain, D. (2007). ‘The Location of Pains,’ *Philosophical Papers*, 36(2): 171–205.
- Bain, D. (2013). ‘What Makes Pains Unpleasant?’, *Philosophical Studies*, 166(1 Supp): 69–89.
- Batty, C. (2010). ‘What the Nose Doesn’t Know: Non-Veridicality and Olfactory Experience,’ *Journal of Consciousness Studies*, 17: 10–17.
- Block, N. (1996). Mental Paint and Mental Latex. In E. Villanueva (Ed.), *Perception*. Atascadero: Ridgeview.

- Block, N. (2006). Bodily Sensations as an Obstacle for Representationalism. In *Pain: New Essays on Its Nature and the Methodology of Its Study*. M. Aydede, Ed. Cambridge, Massachusetts: MIT Press.
- Burge, T. (2010). *Origins of Objectivity*. Oxford: Oxford University Press.
- Byrne, A., and Hilbert, D. R. (1997). Colors and Reflectances. In *Readings on Color (Vol. 1)*, Alex Byrne and David Hilbert (Eds.), Cambridge, Massachusetts: MIT Press.
- Byrne, A., and Hilbert, D. (2003). 'Color realism and color science,' *Behavioral and Brain Sciences*, 26(01): 3–21.
- Byrne, A., and Tye, M. (2006). 'Qualia ain't in the head,' *Noûs*, 40(2): 241–255.
- Chalmers, D. (2004). 'The Representational Character of Experience.' In B. Leiter (Ed.), *The Future for Philosophy*, Oxford University Press, pp. 1–24.
- Chisholm, R. M. (1957). *Perceiving: A Philosophical Study*. Ithaca: Cornell University Press.
- Clark, A. (1996). *Sensory Qualities*. NY: Oxford University Press (Clarendon).
- Clark, A. (2000). *A Theory of Sentience*. New York: Oxford.
- Crane, T. (2000). 'Introspection, intentionality, and the transparency of experience,' *Philosophical Topics*, 28(2): 49–68.
- Cutter, B., & Tye, M. (2011). 'Tracking Representationalism and the Painfulness of Pain,' *Philosophical Issues* 21(1): 90–109.
- Dennett, D. C. (1978). Why You Can't Make a Computer that Feels Pain. In *Brainstorms*, Cambridge, Massachusetts: MIT Press.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, Massachusetts, MIT Press.
- Dretske, F. (1995). *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Dretske, F. (1999). 'The Mind's Awareness of Itself,' *Philosophical Studies*, 95(1–2): 103–124.
- Fazekas, P. (2012). *Tagging the World: Descrying Consciousness in Cognitive Processes*. Ph.D. Dissertation, University of Edinburgh.
- Hardcastle, V. G. (1999). *The Myth of Pain*. Cambridge, Massachusetts: MIT Press.
- Harman, G. (1990). 'The Intrinsic Quality of Experience.' *Philosophical Perspectives* 4: 31–52.
- Hellie, B. (2006). 'Beyond Phenomenal Naiveté,' *Philosopher's Imprint*, 6(2): 1–24.

- Hilbert, D. R. (1987). *Color and Color Perception*. Stanford, CA: CSLI Publications.
- Hill, C. (2006). Ow! The Paradox of Pain. In *Pain: New Essays on Its Nature and the Methodology of Its Study*. M. Aydede, Ed. Cambridge, Massachusetts: MIT Press.
- Hill, C. (2009). *Consciousness*. Cambridge University Press.
- Jackson, F. (1977). *Perception: A Representative Theory*. Cambridge: Cambridge University Press.
- Jackson, F. (2004). 'Representation and Experience.' In *Representation in Mind: New Approaches to Mental Representation*, edited by H. Clapin, P. Staines and P. Slezak, Elsevier.
- Kind, A. (2003). 'What's so transparent about transparency?', *Philosophical Studies*, 115(3): 225–244.
- Kind, A. (2013). The Case Against Representationalism About Moods. In *Current Controversies in the Philosophy of Mind*, Uriah Kriegel, ed., Routledge Press.
- Kulvicki, J. (2004). 'Isomorphism in information carrying systems,' *Pacific Philosophical Quarterly*, 85(4), 380–395.
- Kulvicki, J. (2005). 'Perceptual content, information, and the primary/secondary quality distinction,' *Philosophical Studies*, 122: 103–131.
- Loar, B. (2003). 'Transparent experience and the availability of qualia.' In *Consciousness: New Philosophical Perspectives*, edited by Q. Smith and A. Jokic, Oxford University Press, 77–96.
- Lycan, W. G. (1996). *Consciousness and Experience*. Cambridge, MA: MIT.
- McGinn, C. (1982). *The Character of Mind*. Oxford: Oxford University Press.
- Mole, C. (2010). 'The Content of Olfactory Experience,' *Journal of Consciousness Studies*, 17(11–12): 173–179.
- Moore, G. E. (1903). 'The Refutation of Idealism,' *Mind*, 12: 433–453
- Nida-Rümelin, M. (2007). 'Transparency of Experience and the Perceptual Model of Phenomenal Awareness,' *Philosophical Perspectives*, 21: 429–455.
- O'Sullivan, B., & Schroer, R. (2012). 'Painful Reasons: Representationalism as a Theory of Pain,' *The Philosophical Quarterly*, 62: 737–758.
- Pautz, A. (2010). 'Do Theories of Consciousness Rest On a Mistake?', *Philosophical Issues*, 20(Philosophy of Mind issue): 333–367.
- Peacocke, C. (1992). 'Scenarios, contents and perception.' In Crane, T., Editor, *The Contents of Experience*. Cambridge University Press, Cambridge.

- Pitcher, G. (1970). 'Pain Perception,' *The Philosophical Review*, 79(3): 368–393.
- Reid, T. (1764/1872). *An Inquiry into the Human Mind*. In *The Works of Thomas Reid, D.D.: Now Fully Collected*. Sir William Hamilton, Dugald Stewart (Eds.). Edinburgh: MacLaughlan and Stewart.
- Pylyshyn, Z. W. (2007). *Things and Places*, Cambridge, MA: MIT.
- Shoemaker, S. (1994). 'Self-Knowledge and 'Inner Sense': Lecture III: The Phenomenal Character of Experience,' *Philosophy and Phenomenological Research*, 54(2): 291–314.
- Shoemaker, S. (2000). 'Introspection and Phenomenal Character,' *Philosophical Topics*, 28(2): 247–273.
- Siewert, C. (2004). 'Is experience transparent?', *Philosophical Studies*, 117(1): 15–41.
- Smart, J. J. C. (1959). 'Sensations and brain processes,' *The Philosophical Review*, 68(2), 141–156.
- Smart, J. J. C. (1975). On some criticisms of a physicalist theory of colors. In *Philosophical Aspects of the Mind-Body Problem*, C. Cheng (ed.), Honolulu: University Press of Hawaii.
- Stoljar, D. (2004). 'The argument from diaphanousness.' In *Language, Mind and World: Special Issue of the Canadian Journal of Philosophy*, edited by M. Escurdia, Robert J. Stainton & Christopher D. Viger, University of Alberta Press, 341–90.
- Thompson, B. (2009). 'Senses for senses,' *Australasian Journal of Philosophy*, 87(1): 99–117.
- Tye, M. (1995). *Ten Problems of Consciousness: A representational theory of the phenomenal mind*. Cambridge: MIT.
- Tye, M. (1996a). 'The function of consciousness,' *Noûs*, 1(3): 287–305.
- Tye, M. (1996b). 'Orgasms Again,' *Philosophical Issues*, 7: 51-54.
- Tye, M. (1997). A Representational Theory of Pains and their Phenomenal Character. In *The Nature of Consciousness: Philosophical Debates*. N. Block, O. Flanagan and G. Güzeldere, Eds. Cambridge, Massachusetts: MIT Press.
- Tye, M. (2000). *Consciousness, Color, and Content*, Cambridge, Mass: The MIT Press.
- Tye, M. (2002). 'Representationalism and the Transparency of Experience,' *Noûs*, 36(1): 137–151.

Tye, M. (2006a). Another Look at Representationalism about Pain. In *Pain: New Essays on Its Nature and the Methodology of Its Study*. M. Aydede, Ed., Cambridge, Massachusetts: MIT Press.

Tye, M. (2006b). Reply to Commentaries. In *Pain: New Essays on Its Nature and the Methodology of Its Study*. M. Aydede, Ed., Cambridge, Massachusetts: MIT Press.