Note to readers: those looking for a quick read might skim sections 2 and 3 and then only read 4 and 5, which are the main sections, skipping the rest of the paper. This is a very (very) rough draft – comments, objections very welcome.

The Real Trouble for Phenomenal Externalists:
The Science of Taste, Smell and Pain

Adam Pautz

It is hard to avoid the impression that the biology of the brain is what matters to consciousness—at least the kind we have

-- Ned Block

Abstract: Many defend a version of phenomenal externalism we might call *tracking intentionalism:* David Armstrong, Michael Tye, Fred Dretske, Alex Byrne, and David Hilbert, among others. Christopher Hill and William Lycan defend views in the vicinity which I will also discuss. The usual arguments against tracking intentionalism are based on intuitions about far out cases like Horgan and Tienson's brains in vats, Shoemaker's inverted spectrum, Levine's fish creature, and Block's inverted earth. I will suggest that these *apriori* arguments clearly fall short, despite the great deal of attention they have received. Tracking intentionalism has not been refuted and deserves to be taken very seriously. However, I will develop new *empirical* arguments against it. I show that my arguments are not open to criticisms of empirical arguments against tracking intentionalism raised by Byrne and Tye (2006) or more recent objections raised by Dretske, Lycan and Hill. I also argue that an alternative to tracking intentionalism defended by Uriah Kriegel and Jesse Prinz (and perhaps Gilbert Harman) is not promising. Finally, I briefly develop a puzzle about sensory consciousness, one that has been nearly entirely ignored (aside from some brief remarks by Colin McGinn). I also suggest that much more work needs to be done in order to show that a reductive version of intentionalism can be developed across all the sense-modalities in a way that is in line with the empirical facts.

1 Introduction

Many advocate *intentionalism* about phenomenal consciousness. Roughly speaking, on intentionalism, an episode of sensory consciousness represents or presents sensible properties in external items or in bodily regions, for instance colors, smell qualities, felt pains. And phenomenology is necessarily linked with what sensible properties are represented. Phenomenology and intentionality are inseparable. Intentionalism, in this broad sense, is widely accepted.[1]

Intentionalism leads naturally to *phenomenal externalism:* to a first approximation, the claim that intrinsic duplicates can differ phenomenally. For once intentionalism is accepted, the puzzle of consciousness divides into two puzzles: the old puzzle of where to locate in the physical world the sensible properties, and the newer puzzle of how we manage to represent them. And it is natural to provide a two-part solution to this problem, one which entails phenomenal externalism. First, instead of taking a traditional view that locates sensible properties in the mind, we kick them downstairs into the mind-independent world: colors are reflectance properties, smells are chemical properties of odorants, pains

[1] The disjunctivists/naïve realists reject intentionalism. But I won't be discussing them here, although I do think that this type of view, as well as tracking intentionalism, does not fit well with what science has to tell us about the physical basis of experience (Pautz forthcoming a).

are types of bodily disturbance, and so on. Second, we claim that we manage to sensorily represent these properties – they manage to be disclosed to our minds - thanks to the fact that our sensory systems track them, in the way that a thermometer track temperature. The result is that phenomenal consciousness is just a special kind of information-carrying. Call this kind of view *tracking intentionalism*. It entails phenomenal externalism but tracking intentionalists regard this as an interesting discovery as opposed to a *reductio ad absurdum*: if intrinsic duplicates can suitably track different external properties, then, on this view, intrinsic duplicates can differ phenomenally.

Tracking intentionalism is popular. David Armstrong is the founding father not only of intentionalism but also of tracking intentionalism (although he often does not get the credit he deserves).[2] He had a sophisticated psychosemantics long before the term was coined. More recently, tracking intentionalism has been taken up by Fred Dretske, Michael Tye, and others. I think Alex Byrne and David Hilbert and Christopher Hill are also committed to something like tracking intentionalism, but their views are somewhat less developed (see footnote 20).[3] Tracking intentionalism is not just popular; it's also important. Many think intentionalism about consciousness of some form has to be true. And it may be that our best reductive theories of sensory intentionality are broadly speaking tracking theories. If these points are correct, then something in the vicinity of tracking intentionalism might represent our best chance at reducing consciousness. You might think that tracking theories are too simple and that there are better, more sophisticated reductive versions of intentionalism out there. I disagree: I understand "tracking theories" very broadly (see the end of §4 and note 20) and I think that if we have to reject such theories (as I will argue we must) intentionalists are kind of up the creek when it comes to reducing consciousness and its intentionality. So the stakes are high.

Yet I will argue that tracking intentionalism is mistaken. My arguments against tracking intentionalism will have two unique features.

First, the usual arguments against externalist views like tracking intentionalism are based on intuitions about extreme cases like Horgan's brains in vats, Shoemaker's inverted spectrum, Speaks psychedelic phenomenology, Levine's fish creature, and Block's inverted earth. But, for reasons I will briefly explain (in §2), I think these arguments are unsuccessful. At least some forms of tracking intentionalism cannot be ruled out on the basis of intuition. Instead my main arguments are firmly rooted decades of research in neuroscience and psychophysics. I think these arguments, together some apriori arguments (to be discussed in §6) *not* in the literature, are the *only* good arguments against tracking intentionalism. As we shall see, the arguments are not open to the objections Byrne and Tye

[2] Armstrong didn't formulate his form of intentionalism in terms of supervenience, as it is typically formulated today; but he advocated a reductive version of intentionalism which clearly entails the supervenience of phenomenal character on intentional content. He also early on explained experience in terms of belief, which contemporary intentionalists would not in general do (but see Byrne 2008 and Jackson 2004). But later (1984) Armstrong backed off of the belief formulation and preferred to explain experience in terms of 'information carrying' states.

[3] It might seem that Byrne and Hilbert are not committed to tracking intentionalism about taste, smell, or pain. They express skepticism about current reductive theories of intentionality (2003); and they do not discuss taste and smell, mainly focusing on color. But I will argue that they are committed to tracking intentionalism, in a sense to be explained (see footnote 20).

(2006) have raised against the attempt to undermine tracking intentionalism on empirical grounds.

Second, whereas most arguments focus on visual experiences, mine focus on non-visual sense modalities, in particular taste, smell and pain. Tracking intentionalists have mostly discussed color and color experience. In my opinion, even here they have not adequately developed the view. But they have said even less about the chemical senses and pain. They do not bother to get down to the nitty-gritty. How can they be confidant reductionists if they have not attempted to deliver the goods? It is true that Clare Batty and Casey O'Callaghan have attempted reductive accounts of smell qualities and sound qualities (respectively) but they have not attempted reductive accounts of our *experience* of these qualities. If they accept something in the vicinity of tracking intentionalism, they faces the problems I will be raising.[4]

So my main target is tracking intentionalism. But I will also briefly argue against a version of phenomenal externalism in the vicinity of tracking intentionalism that is defended by Uriah Kriegel and Jesse Prinz (and perhaps Gilbert Harman).

To be clear, my target is *not* phenomenal externalism, the general thesis that intrinsic duplicates can differ phenomenally. Indeed, while I reject tracking intentionalism, I think that other forms of phenomenal externalism might well be true. As I said, I do not think phenomenal externalism can be ruled out *apriori*. So I won't here be defending internalism about phenomenology over externalism. My aim is just to argue against tracking intentionalism, a view which the leading phenomenal externalists defend and which I think might represent the intentionalist's best shot at reducing phenomenal consciousness. It is in this sense that I will be raising troubles for the phenomenal externalists.

My main aim, then, will be the negative of showing tracking intentionalism does not work. But there will also be a positive upshot. By seeing how tracking intentionalism goes wrong, we will see what the right theory of sensory consciousness might look like.

My plan is as follows. First (§2) I say what tracking intentionalism is and briefly say why I think it cannot be ruled out apriori. Then (§3) I provide some very basic empirical background. Next (§§4-6) I explain my three arguments against tracking intentionalism. Then (§7) I reply to objections raised by Tye, Byrne, Hill, Lycan and Dretske. Next (§8) I briefly argue against a theory in the vicinity of tracking intentionalism that is defended by Uriah Kriegel and Jesse Prinz (and perhaps Gilbert Harman). Finally (§9) I'll suggest we need a radically different approach, and briefly describe a puzzle about consciousness that naturally arises out of the discussion.

2 What Is Tracking Intentionalism?

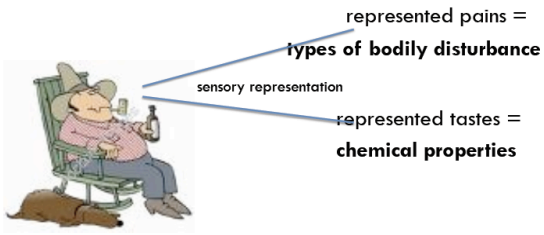
To begin with, let me say in great detail what tracking intentionalism is. Many defend tracking intentionalism but I will focus on the views of Tye and Dretske. So, for the present, tracking intentionalism can just be taken to be the disjunction of their views. A better, more general formulation will be provided in §4. I will

[4] I do not discuss sound here. But the sorts of problems I will discuss here carry over to tracking intentionalism about sounds and sound experiences, as I explain elsewhere: https://webspace.utexas.edu/arp424/www/rest.pdf

also argue (note 20) that Byrne and Hilbert and Hill are apparently committed to tracking intentionalism in this more general sense.

In the cases of taste, smell and pain that will be my focus here, Tye and Dretske's tracking intentionalism has following three parts.

The first part is intentionalism. On a simple version of intentionalism, having a pain experience with a certain phenomenal character is a matter of standing in a relation, which I will call the *sensory representation relation*, to a pain-property ostensibly located in a certain bodily region. The phenomenal character of the experience is fully determined by the property represented. In cases of phantom limb and the like, the represented pain property is not instantiated where it seems to be, so there is non-veridicality. Likewise, having a taste experience with a certain phenomenal character is a matter of standing in the sensory representation relation to a taste-property ostensibly located in or around the tongue. The story for smell experience is similar.



The second part of tracking intentionalism is a response-independent and reductive theory of the represented sensible properties. (I count views like Chris Hill's in his book *Consciousness* and in other recent writings on which sensible properties are viewpoint relative properties as response-independent views, provided that such properties are not explained in terms of relations to subjects.) What does this mean? Tracking intentionalists deny that the sensible properties are Lockean dispositions to produce experiences. Instead they hold that they are physical properties of external items. What physical properties? Tracking intentionalists have not fully addressed this question. They have discussed colors but have said very little about other sensible properties. One idea would be that felt pain qualities, tastes qualities, and smell qualities are properties involving peripheral neural activations in the body, tongue and nose. We are going to look at this idea later (§7) to see if it might get the tracking intentionalists out of the empirical problems I will be raising – I will suggest it does not. But in fact that is not what they say. Instead they identify the sensible properties with the external physical properties that set up the peripheral neural activity. Armstrong said that taste and smell qualities are "[physical] qualities of objective physical phenomena" such as "motions of molecules" without saying anything more definite. Tye and Dretske follow suite but say little more. As for pains, Tye and Dretske say that they are different types of bodily damage.[5] A problem here is that many pains don't involve damage. Christopher Hill says that in these cases the felts pains are different kinds of potential damage. But: what kinds of properties are these? Dispositions-to-be-damaged? How are they represented, if representation requires causation and potentialities don't cause? And what about pains that don't even

[5] See Armstrong 1984, Tye 1995, and Dretske 2003.

involve *potential* damage? So tracking intentionalism is very under-developed, even though a lot hangs on it because it might represent our best chance at reducing sensory consciousness. However, the view is clear enough for our purposes. The idea is that smell and taste qualities are chemical properties of foods or odor clouds and felt pain qualities are bodily disturbance types. Whereas traditional views of consciousness locate the sensible properties in the mind, *tracking intentionalists kick them downstairs into the mind-independent world*.

The third and last part of tracking intentionalism is a broadly tracking theory of how we represent these properties. Tye reduces the sensory representation relation to the *optimal tracking relation*, that is, the relation: individual *x* is in an internal state that plays consciousness-making functional role *F* and that, under optimal conditions, would be caused the instantiation of property *Y* (or for short, that is *optimally caused by Y*).[6] Dretske reduces the sensory representation relation to the *indication relation*: individual *x* is in an internal state that plays functional role F and that has the function of indicating *Y*.

The minor differences between Tye and Dretske's views won't matter here, and I will focus mainly on Tye's optimal tracking version. The basic idea is simple: sensible properties are external physical properties, and they become disclosed to the mind by virtue of the fact that our sensory systems carry information or misinformation about them. Phenomenal character is fully determined by what sensible properties become thus present to the mind.

As I said above, tracking intentionalism is important. Many think intentionalism about consciousness of some form has to be true. And our best reductive theories of intentionality in general are externalist. If these points are correct, then something in the vicinity of tracking intentionalism might represent our best chance at reducing consciousness.

Yet some might think that tracking intentionalism is undermined by arguments already in the literature. They might think the view just shouldn't be taken seriously; so the empirical arguments I am going to give are unnecessary. I disagree. In fact, I think that the empirical arguments I am going to give (together with a couple of apriori arguments not in the literature) are the *only* good arguments against tracking intentionalism. Let me briefly review some of the main arguments against tracking intentionalism and indicate why I find them lacking.

(i) Some think that it is undermined by intuitions about far out cases, such as Horgan's brains in vats, Shoemaker's inverted spectrum, Levine's fish creature, Block's inverted earth, Swampman, and so on. These arguments have received an enormous amount of attention but I think tracking intentionalists are right when they say they show nothing. Shoemaker assumes that one individual might have a reddish experience and another have a greenish experience, even if the experiences optimally track the same reflectance property. But while this case is conceivable, the tracking intentionalist will just say it is not possible. This is how Tye

[6] On Tye's PANIC theory, the relevant functional role is, very roughly: being poised to influence the formation of beliefs and desires. Dretske takes a somewhat similar view. Despite the fact that he uses different terminology ('attention') than Tye, Jesse Prinz's AIR theory is ultimately very similar to Tye's PANIC theory. (Prinz himself (2010, 323) thinks his talk of 'attention' makes his theory importantly different from Tye's. I am not sure about this: it only matters if it means that the AIR theory clearly delivers different verdicts than the PANIC theory in some cases, but this has not been shown.) This yields a first-order theory of consciousness. Tracking intentionalism could also be combined with a second-order thought theory (Rosenthal). My arguments will be independent of this issue.

responds.[7] Physicalists of all people should realize conceivability does not entail possibility: after all Zombies are conceivable too. Likewise, Levine (2003; also *Purple Haze*) appears to assume that a fish-like creature could have reddish experience of an object with one eye and greenish experiences with another eye (a split visual field perhaps), even though the experiences optimally track the same reflectance property: spectrum inversion within an individual, so to speak. But, again, why suppose that this is possible? (Lycan 2006 grants that the case is possible; I'm suggesting that he is being needlessly concessive.) It is conceivable but physicalists must admit that this is not enough to prove possibility. Levine does say that the underlying neural states are different. But maybe it is a case of alternative neural realizations of the same experience. We are familiar with such cases across creatures; why couldn't it occur within a creature? There is more to say about Levine's case[8] but one might worry that to the tracking intentionalist it will be no more convincing than the familiar case of spectrum inversion.

Block and Horgan and others think that internal duplicates – for instance, me and my twin on inverted earth, or me and my brain in a vat duplicate – cannot differ phenomenally. Why think this? One idea is that we can know apriori that intrinsic duplicates can't differ phenomenally. This seems totally wrong to me. A

[7] See Tye 2000, chapter 5.

[8] Levine's argument (2003, also in *Purple Haze*) against tracking intentionalism actually has an epistemological element that I have not discussed. The idea seems to be that on tracking intentionalism the following should be metaphysically possible: a fish (or a human) creature could, on the basis of seeming to have two radically different experiences, believe that it has radically different experiences, viz. reddish and greenish experiences, via different eyes; and then the creature could discover that the two experiences track the same reflectance and (provided God or someone also told him tracking intentionalism is true!) thereby discover that it is really having one color experience. But, the argument continues, we should not say this is metaphysically possible. As Levine puts it, 'this doesn't make sense'. In my view, this argument has serious problems. (i) Again – and I think this is perhaps the best response – the tracking intentionalism might say that on his view there is no reason to believe that this situation is metaphysically possible and reason to believe it is not. On his view, if the two states track the same reflectance, then they realize the same color experience (a case of multiple realization of the same experience within a creature). (Or, on *some* versions of tracking intentionalism, maybe they realize no experience: cf Speaks point in footnote 10 of the present paper.) And, he might say, if the creature has one experience (e. g. a reddish experience), it is metaphysically impossible that it should seem to itself to have two radically different color experiences (a reddish and a greenish experience), and believe that it has two such radically different color experiences. (Many say that there cannot be radical divergences between how our experiences in fact are, on the one hand, and how they seem to one, and how we believe them to be, on the other hand; there is no reason why the tracking intentionalist cannot say this too.) So, the tracking intentionalist might say that from his point of view the conditions present in Levine's (introspectively seeming to have two radically different experiences while tracking the same reflectance) case are inconsistent. (ii) I would also point out that Levine's argument, if it were successful, would generalize in a bad way against every theory of experience. For instance, consider the identity theory. If Levine's argument against tracking intentionalism is good, one can construct a similar argument against this theory. By Levine's reasoning, if identity theory is true, then the following situation is possible: a fish (or a human) creature could, on the basis of seeming to have two radically different experiences, believe that it has radically different experiences, viz. reddish and greenish experiences, via different eyes; and then the creature could discover that the two experiences are constituted by the very same brain state and (provided God or someone also told him tracking identity theory is true!) thereby discover that it is really having one color experience. So, given that Levine is right that this situation makes no sense, we could rule out the identity theory as well as tracking intentionalism – and, mutatis mutandis, any theory of experience (including dualist theories!). Of course, this does not seem right. So there must be something wrong with the argument. I have already identified the flaw: the theorist can always say that his theory does not have the consequence that the relevant scenario is possible, as I explained in connection with tracking intentionalism under (i).

very natural way of thinking about experience says that at least in good cases it gets it character from the presented or represented character of external world. Both naïve realists and intentionalists try to accommodate this way of thinking of experience. Now why couldn't it be that we could hooked up to that world by what of extrinsic relations involving tracking? How else could it happen? Then duplicate brains might differ phenomenally, provided they get hooked up to different environments, in which different sensible properties are instantiated. Or, as in the case of the brain in the vat, provided one of the brains fails to make actual or potential contact with the world at all. On this view, the brain plays an enabling role; it is not all by itself sufficient for experience. This idea does not seem incoherent to me at all. On the contrary, it seems like a quite natural view. Indeed, Euclid and Ptolemy advocated the emission theory of vision, according to you see an external state, and hence have an experience with a certain phenomenal character, when a ray emitted from the eye hits the object. This view seems externalist because presumably duplicate brains, in different environments, might emit noetic rays that hit objects with quite different sensible properties. This view raises obvious problems (what about hallucination and illusion?) and we know it is false (the rays go from the object to the eye, not from the eye to the object) but apparently it cannot be ruled out apriori: very smart folks who grasped the relevant concepts believed it. Indeed a recent study found recent evidence that as many as 50% of American college students believe in emission theory! So, for these and other[9] reasons, I think we can't rule out phenomenal externalism *apriori*.

Of course, it might be thought that we cannot quite rule out phenomenal externalism apriori, we can rule it out on the basis of modest empirical facts about the physical basis of experience. Thus, in addition to relying on apriori intuition, Horgan says that internalism is proven by the fact that the world causes experiences by causing brain states. But this argument is fallacious. For this is equally true of natural-kind beliefs, for which externalism is true. As we shall see, I think relying on empirical considerations is on the right track; but to undermine tracking intentionalism we need more sophisticated empirical arguments, ones which take a more in depth look at what role the brain plays in determining phenomenal character. We will also see that it is possible to refute tracking intentionalism without relying on anything so strong as total internalism: it is enough to establish the weaker thesis that internal factors play a role, in a sense to be explained.[10]

---

[9] See Pautz forthcoming b.

[10] Speaks (unpublished online paper) has an interesting argument against a *specific, covariational version* of tracking intentionalism but at a crucial point it assumes without detailed argument some kind of internalist thesis that the tracking intentionalist would reject. He has us consider a situation in which "S*RED* and S*GREEN* [the brain states realizing our actual experiences of red and green] indicate a single color — say, redness. This would then be a case, if the covariational theory were true, in which neither S*RED* nor S*GREEN* would covary with a color in optimal conditions, and hence would be a case in which neither represents a color." Then he says that on tracking intentionalism in this case if a subject had the brain state Sred he would have 'red phenomenology' but it would not represent redness – which Speaks thinks is not possible because he accepts a kind of intentionalism. But of course the tracking intentionalist will say that this case is just not possible too. He will say that when a subject has the brain state Sred it doesn't represent anything (granting Speaks' point) so he doesn't have any experience at all – much less a reddish one. Indeed, if in this scenario there is no state that track redness, *no one* in this scenario could have a reddish experience, on tracking intentionalism. Speaks seems to be assuming that the Sred brain state realizes a reddish experience even in this scenario in which it fails to covary with redness and hence (on tracking

(ii) Others think that tracking intentionalism is undermined by certain actual cases of variation. They have mostly focused on color vision. For instance, a color chip might look pure blue to John and green-blue to Jane. This is usually discussed as a problem for "response-independent" views on color but one might also think it is a problem for tracking intentionalism. But I think this is not a problem. Some tracking intentionalists take the inegalitarian view that one is right and the other is wrong. I think this is a mistake, for reasons I have explained elsewhere.[11] They ought to take the pluralist view that the chip has both colors; they are overlapping but distinct disjunctions of reflectances; John tracks and so experiences the one and Jane tracks and so experiences the other. The claim that the chip has both colors is somewhat strange but it is not incoherent.[12] Of course, there are many cases like the John-Jane case. Ned Block (2005) has noted that the same bodily disturbance-type can give different individuals slightly different pains. The tracking intentionalist can easily handle this case in the same "pluralist" fashion.

I do think that there are more difficult cases involving other sense modalities but they have not been addressed in the literature; yet in the end I think they too do not bring down tracking intentionalism. For instance, 50% of humans cannot smell the compound androstenone, which is found in armpit sweat and pork. Of course, that's not a problem for tracking intentionalists. The problem comes from the fact that, among 50% who can smell it, half perceive it as having a pleasant sweet floral smell and the other half smell it as having an unpleasant ruinous odor (Keller et al. 2007). Neither group is a victim of any kind of interference or malfunction involved: conditions are optimal.

How might the tracking intentionalist handle this case? The case seems to present a difficulty for the tracking intentionalist. (And, in general, anyone who, like Clare Batty, attempts a response-independent theory of smell.) The tracking intentionalist could try pluralism again. To illustrate, suppose on being presented with an odor cloud containing androstenone John gets the sweet smell and Jane gets the ruinous smell. Let the realizing brain states be *S* and *U*. The idea would be this: John's experience-vehicle *S* is optimally caused by a range *D1* of chemical compounds that includes androstenone. Jane's experience-vehicle *U* is optimally caused by a distinct range *D2* of chemical compounds that also includes antrodostone. (I say a "range" because many different compounds can yield the very same brain state because of "metamerism" and because of variability in the subjects' internal wiring.) Thus sweet smell quality that John represents is *D1* and the ruinous smell quality that Jane represents is *D2*.

---

intentionalism) fails to represent redness – *but it is precisely this which the tracking intentionalist rejects*. To get a convincing argument against tracking intentionalist, we must look seriously at the neuroscience and psychophysics, which as we will see actually provides a non-question begging argument for thinking that tracking intentionalism fails.

[11] Pautz 2010.

[12] One might think it is incoherent because if an object is pure blue it is not at all greenish, so it cannot also be greenish blue. But on one natural view there is no incoherence. On this view, to say that an object is pure blue is to say that it has one color *quality* that is bluish and not at all greenish (etc.); to say that it is greenish blue is to say that it has a color quality that is greenish and bluish. The same object can have one quality that is not at all greenish and *another* quality that is greenish. Since two qualities are involved, and since it is qualities and not object that in the first instance have hue qualities, we avoid the result that one entity is both not greenish and greenish.

The trouble with this is that the odor cloud containing androstenone would now possess both *D1* and *D2*. So it would objectively be both sweet-floral and ruinous. It's hard enough to suppose a color chip is both pure blue and a bit greenish-blue; it is even hard to suppose that an odor cloud might "objectively" have two such radically different smells.[13]

But I suppose that the tracking intentionalist could just bite the bullet. He might say we should be naturalists and shouldn't put too much stock in our "intuitions" – which after all have been overturned in the past. The overall best theory of experience might have some surprising consequences.

(iii) Finally, there is an argument from failure. Part of tracking intentionalism is a reductive account of sensory intentionality. But the project of reducing intentionality kind of dead-ended in the 90s, and many despair of ever completing the project. The causes for despair are the many well-known problems of detail: the depth (distance) problem, the disjunction problem, the problem of co-extensive properties, the indeterminacy problem, and so on. Mark Johnston (2007) has recently inveighed against the whole reductive program on the basis of a number of considerations. But in response many tracking intentionalists like Tye, Byrne and Lycan have simply expressed a "fingers-crossed" attitude: they say it's still early days, and that these problems have some naturalistic solutions, even if we have not (and maybe cannot) come up with their solution. They might also point out that sensory intentionality is simpler and more primitive than cognitive intentionality, so perhaps it is more amenable to reductive explanation. It's hard to know how to respond.

In short, there is reason to believe that existing arguments against tracking intentionalism fall short. I believe that my arguments against tracking intentionalism succeed where others fails. They do not rely on dubious intuitions, but on empirical research. Further, they do not rely on problems of detail. They show that, even if we wish away the problems of detail, tracking intentionalism is a radically incorrect theory sensory consciousness.

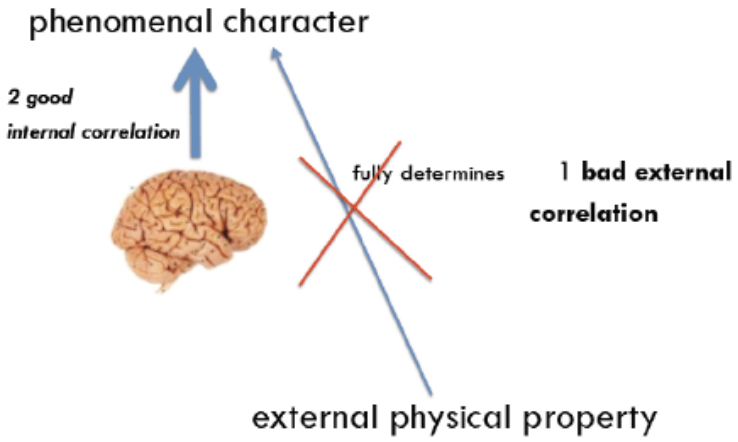
3 Empirical Background

As I said, my arguments against tracking intentionalism are empirically-based. So let me begin with some *very basic* empirical background. In this section, I am just going to report the empirical facts; I'm not trying to argue for "internalism" and against tracking intentionalism. The arguments against tracking intentionalism will come later.

The three parts of tracking intentionalism entail that phenomenal character of an experience is fully determined by the external physical property that is optimally tracked, and thereby represented, by the experience. But I believe that decades of research in psychophysics and neuroscience undermine this externalist picture. It is a case of a profound disconnect between philosophy and science. In fact, in view of this long-standing research, it is really surprising that tracking intentionalism has received the level of attention it has from philosophers.

In particular two empirical facts are relevant. They are illustrated here:

[13] Another case, better known to philosophers, is that of phenyl – PTC. It is tasteless to about 30% of humans and bitter to the rest. This is not a problem for tracking intentionalists: they can say that the tasters correctly represent it as bitter whereas the others just fail to represent its bitterness.



First, for decades psychophysics has revealed that, even under optimal conditions, there is an extremely bad correlation between phenomenal character and the character of the external physical property tracked, the property which tracking intentionalists claim to determine phenomenal character. Roughly, what that means is that structural relations among experiences (similarity and difference, equal intervals, proportion) are not matched by structural relations among the external physical properties optimally tracked. For short, there is *bad external correlation*. This is quite different, by the way, from "metamerism".[14]

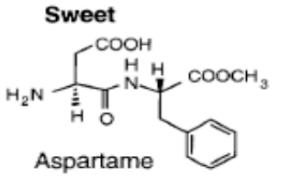
Second, neuroscience has revealed that phenomenal character is much better correlated with neural firing patterns in the brain. What I mean goes beyond the mere claim that experiences have neural correlates. Of course they do. This does nothing to undermine tracking intentionalism. (Experiences are also correlated with properties of the form *optimally tracking external property F* – the properties that the tracking intentionalist identifies experiences with.) What I mean is that structural relations among experiences (similarity and difference, equal intervals, proportion) are well matched by structural relations among their neural correlates. For short, while there is bad external correlation, there is *good internal correlation*. As we shall see in subsequent sections, this does provide the materials for empirical arguments against tracking intentionalism.

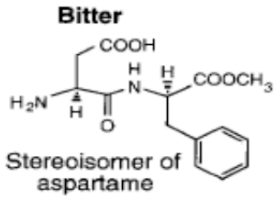
To illustrate, I will first discuss the chemical senses, taste and smell; then I will turn to pain.

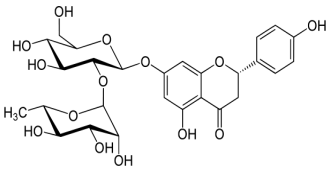
3.1 Taste

[14] "Metamerism" is familiar from the science of color vision. But the same phenomena is present in all sense modalities. Thus, for instance, a variety of different combinations of molecularly different odorants, and different levels of concentration, can produce the very same smell experience, S. This is easily handled by tracking intentionalists: they can say that S represents a disjunctive chemical property, D. The kind of "bad external correlation" I will discuss in what follows is quite different. Let S1, S2 and S3 be smell experiences and let D1, D2 and D3 be the disjunctive chemical properties which, on tracking intentionalism, they represent and which determine those experiences' phenomenal characters. An example of what I mean by bad external correlation is this: it might be that S1 is more like S2 than S3, even though it is not the case that D1 is more like D2 than D3. Whereas "metamerism" is easily handled by tracking intentionalism, "bad external correlation" together with "good internal correlation" undermine the view, as we shall see in later sections.

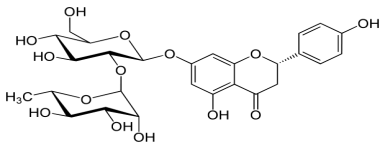
The tracking intentionalist will presumably say that different types of taste qualities are different types of chemical properties that are tracked and thereby represented by our taste experiences. The phenomenal character of a taste experience is fully determined by the response-independent character of the chemical type it optimally tracks. But, in the case of taste, examples of bad external correlation abound. Here is one example:







Strychnine



Naringnin

Suppose you first taste the chemicals represented in the bottom row and then those represented in the top row. Then, as a matter of fact, your first two taste experiences will resemble each other more (similar bitter taste experiences) than your second two taste experiences (one sweet and the other bitter). But evidently the chemical properties tracked (and hence on tracking intentionalist represented) by your first two experiences (viz. those in the bottom row) do not in any sense resemble more than those tracked by your second two taste experiences (viz. those depicted in the top row). This is an example of what I mean by bad external correlation. Indeed, the chemicals depicted in the top row are stereoisomers of aspartame: chemicals that have the same molecular formula and sequence of bonded atoms, and differ (sometimes only slightly) in the three-dimensional orientations of their atoms in space (Walters 1996).

In general, resemblances and differences among taste quality are not matched by resemblances and differences among chemical properties. Two tastes might be phenomenally very different, but optimally track chemical properties with very similar physical structures. For instance, trehalose has a distinctly sweet taste, but gentiobiose is bitter, even though they are very similar disaccharides composed of two glucose units. Neohesperidin is a bitter compound found in citrus peel; removing a single carbon-oxygen bond produces the intensely sweet-tasting compound neohesperidin dihydrochalcone. In fact, as we just saw, even very similar *stereoisomers* can taste radically different. Likewise, while as mentioned gentiobiose tastes bitter, it has an anomer (a kind of very similar stereoisomer), namely isomaltose, that tastes sweet (Sakurai 2010).

There are just as many examples of phenomenal similarity despite radical chemical differences, especially when it comes to bitter-tasting compounds. Bitter-tasting compounds form a very heterogeneous lot that includes moderately

large organic compounds such as the citrus compound naringin, the large organic acids found in hop oils, small molecules like urea, and even (as we saw above) some sugars. Van der Heijden (1993) listed no fewer than 19 distinct chemical families of bitter substances. Despite being very physically different, these compounds can optimally produce in us very similar bitter experiences (Wolfe 2009). This make evolutions sense: they are all bad for us, so the body has no need to distinguish them.

So psychophysics has shown that, in the case of taste, there is bad "external" correlation between the phenomenal and the chemical. By contrast, neuroscience has shown that there is a good "internal" correlation between the phenomenal and the neural.

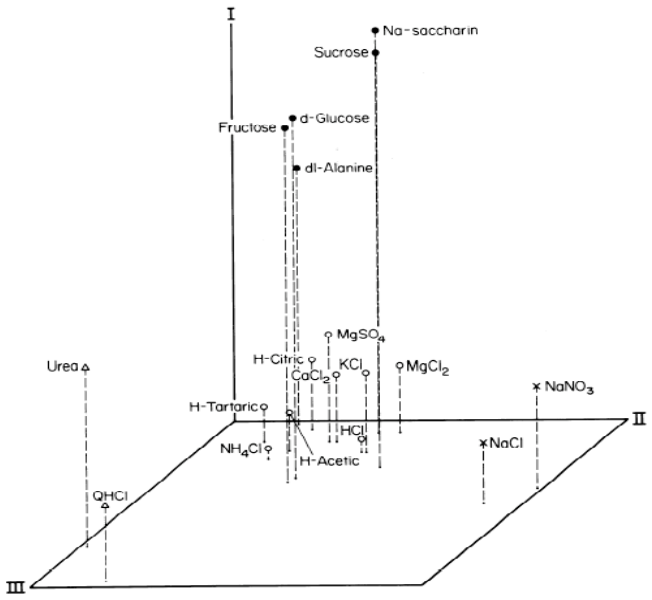
On the tongue there are several types of taste sensitive receptors, each optimally responsive to substances that we regard as having one of the "basic tastes" (sweet, salty, bitter, sour, unami). But at more central locations in the taste system neurons are tuned broadly, with many neurons responding to more than one taste quality. So when one experiences a particular taste one undergoes a distinctive pattern of neuronal firing across many centrally-located neurons. This is called an *across-fiber pattern* or *ensemble activation*.

The resemblance ordering among tastes is well-matched by the resemblance-ordering among such ensemble activations (as determined by multidimensional scaling). This is something that Paul Churclhand and Austen Clark have emphasized in arguing for a brain-based approach to sensory consciousness.[15] For instance, Erickson (1963) shocked rats while they drinking a solution containing potassium chloride, so that they came to avoid it. Afterwards, they also avoided solutions containing ammonium chloride but they did not avoid solutions containing sodium chloride. This suggests that the rat's experience of ammonium chloride is much more similar to its experience of potassium chloride than its experience of sodium chloride. Erickson also found that the *ensemble activation* produced in the rat's chorda tympani by ammonium chloride is much more similar to that produced by potassium chloride than that produced by sodium chloride. This is a pattern of neural firing across a number of neurons. Further, when Schiffman and Erickson (1971) asked humans to make similarity judgments between a number of different solutions, they found that similarities and differences in quality corresponded remarkably well to similarities and differences in ensemble activations in the rat: in the main, the greater the phenomenal similarity, the greater the ensemble activation similarity.

In the Figure below a multidimensional space of ensemble activations in the hamster taste system is shown. Distances among points represent degree of similarity. The space is clearly interpretable on the basis of human taste experiences. The ensemble activations of sweet tasting-substances (sucrose, fructose, d-glucose, Na saccharin, and Galanine) are very similar to each other and very different from those of bitter-tasting substance (QHCl and urea). Likewise the ensemble activations of sour-tasting and salt-tasting substance cluster together in different areas. [16]

[15] See for instance Churchland *The Engine of the Soul* and Clark *Sensory Qualities*.

[16] In the case of taste (as well as in the case of smell) temporal coding might play an important role. This is an area of current interest. Cf Di Lorenzo 2009 and Laurent 1996.



From Smith 1983

(So far we have focused on the neural correlates of taste quality. There is also a very good correlation between average taste-cell firing rates and *taste intensity*. Due to an anatomical peculiarity, the chorda tympani nerve can be accessed in humans during middle-ear surgery by means of an electrode. In a well-known experiment, Borg et al. (1967) had patients estimate numerically taste magnitudes of certain substances at various concentrations. Then they recorded from taste cells and found nearly perfect agreement between the neural data.)

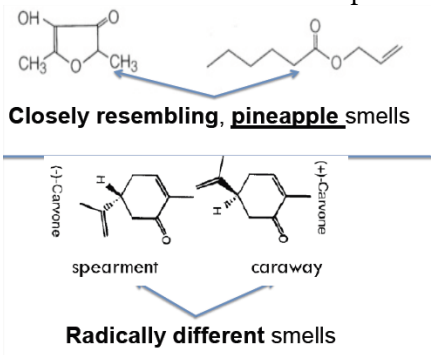
3.2 Smell

Now let us turn to smell, which, as is well known, informs of our experience of taste. As in the case of taste, the tracking intentionalist will presumably say that different types of smell qualities are different types of chemical properties that are tracked and thereby represented by our smell experiences. The phenomenal character of a smell experience is fully determined by the response-independent character of the chemical type it optimally tracks. But, in the case of smell, examples of bad correlation between phenomenal character and the character of the physical properties optimally tracked are even more plentiful than in the case of taste (Beets, 1970; Polak, 1973; Arctander, 1969). Cowart and Rawson (2005) sum up the situation as follows:

> Available evidence indicates that numerous chemical and molecular features (e.g., molecular weight, molecular mass and shape, polarity, resonance structure, types of bonds and sidegroups) can all influence the odorous characteristics of a chemical. However, no systematic description of how these characteristics relate to particular odor qualities has been developed. Odor qualities have been classified in a variety of ways, but schemes based on perceptual qualities reflect neither chemical nor physiological referents. In other words, chemicals that bear little resemblance structurally can smell the same, and

chemicals that are nearly identical structurally can elicit very different perceptual qualities.

Here I will focus on one example:



From Goldstein Sensation and Perception and Linster 2001

The first two (top two) molecules are clearly more different from each other than the second two. Indeed, the second two molecules are enantiomers of carvone (-, +): stereoisomers that are *mirror images*. Yet, as a matter of empirical fact, one's experiences of the first two molecules are more similar to each other than one's experiences of the second two. In particular, the first two produce very similar (but non-identical) pineapple smells, whereas the second two smell like spearmint and caraway, respectively. So there is "bad correlation" between the resemblances among the experiences and the resemblances among the chemical properties which those experiences track and hence (on Tye and Dretske's views) represent. Yet on tracking intentionalism it is the chemical properties represented are supposed to be what fully determines the experiences' phenomenal character!

Another interesting psychophysical fact is worth mentioning. The tracking intentionalist would presumably say that the intensity of a smell experience is constituted by the level of concentration it represents. This makes senses, because in general changes in the level of concentration go with changes in intensity. But sometimes a mere change in the concentration of a chemical can strikingly alter the quality of olfactory experience: for instance, the smell experience of *thioterpineol* is described as "tropical fruit" at a low concentration, as "grapefruit" at a higher concentration, and as "stench" at a still higher concentration (Malnic et al. 1999). Here the radically externalist account of phenomenal character promoted by tracking intentionalists is at an explanatory disadvantage. Why should changes in represented concentration sometimes constitutes changes in intensity, sometimes changes in quality? As we shall see, the puzzle is resolved by a more internalist view on which internal neural factors play a role in determining phenomenal character.

In the case of smell as in the case of taste, while psychophysics has revealed very bad external correlation, neuroscience has shown good internal correlation.

Humans have about 450 *types* of smell receptors on the olfactory epithelium of the nose. (Contrast this with the mere three cone-types in vision or the mere four or dive receptor types for taste.) They synapse at the olfactory bulbs, which in turn are connected to the primary olfactory cortex. The primary olfactory cortex is subdivided into several different areas: the anterior olfactory cortex, the ol-

factory tubercle, the piriform cortex (about which more presently), parts of the amygdala and the entorhinal.

It is well known that chemicals that are nearly identical structurally can elicit very different smell experiences. I have focused on the case of the enantiomers of carvone but there are many such cases. Malnic and coworkers (1999) found that in such cases the very similar chemical produce very different patterns of firing across the smell receptors in mice. So where there is bad external correlation there is good internal correlation. Interestingly, Malnic et al also found that at different concentrations the same chemical can sometimes produce radically different *patterns* of activation across the smell receptors. This goes with the fact – just mentioned above as a puzzle on an external view of phenomenal character – that a mere change in the concentration of a chemical can sometimes strikingly alter quality of olfactory experience, not just the intensity.

We also saw that even enantiomers (chemicals that are mirror image rotations) can smell quite different. The explanation is that the rotated molecules don't fit the same smell receptors (as if you were trying to fit your right hand into your left hand glove). Because of this, they stimulate different receptors. This is because the receptors contain chiral groups, allowing them to respond more strongly to one enantiomer than to the other. Consistently with this, Linster and coworkers (2001) found that enantiomers that smell quite different (as determined by behavioral measures) such as those shown in Figure 4b also produce quite different neural patterns further downstream in the olfactory bulb of rats. Presumably it is the same in humans.[17]

But probably the most striking demonstration of "good internal correlation" in the case of smell is provided by a recent fMRI experiment by Howard and coworkers on *human subjects* (2009). This experiment is an advance in several ways. Most obviously, by contrast to animals, human subjects are able to *report on* phenomenal similarities and difference in their smell experiences. By obtaining their reports, and by performing fMRI scans, they obtained very strong evidence of "good internal correlation" in the case of smell.

In their main experiment, for a reason that will emerge, Howard et al. used chemicals that are physically very different but smell similar (viz. minty, woody or citrus); in other words, they focused on cases where there is "bad external relation": see below.

[17] Linster et al. also found cases where two enantiomers smell the same; in such cases they produced nearly identical olfactory bulb patterns. There are several elegant studies which have demonstrated that odor-evoked spatial activity in the rodent olfactory bulb correlates with behavioral measures of odor similarity (cited by Howard et al. 2009).

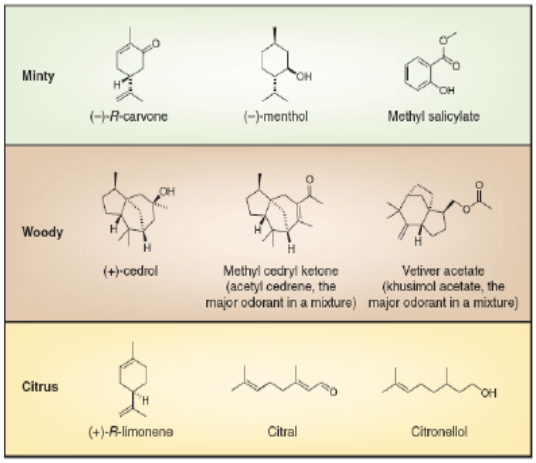


Figure 5: the odorants used by Howard et. al. 2009. From Margot 2009.

Using these odorants, Howard et al. found "that spatially distributed ensemble activity in human posterior piriform cortex (PPC) coincides with perceptual ratings of odor quality, such that odorants with more (or less) similar fMRI patterns were perceived as more (or less) alike" (2009, 932). (See Figure 6.) In particular, Howard and coworkers found that, even though the molecular structures in each of the three families are quite different (shown in the horizontal rows in Figure 5), they produce very similar ensemble activations in PPC, which are distinct from the activation patterns of the other two categories.

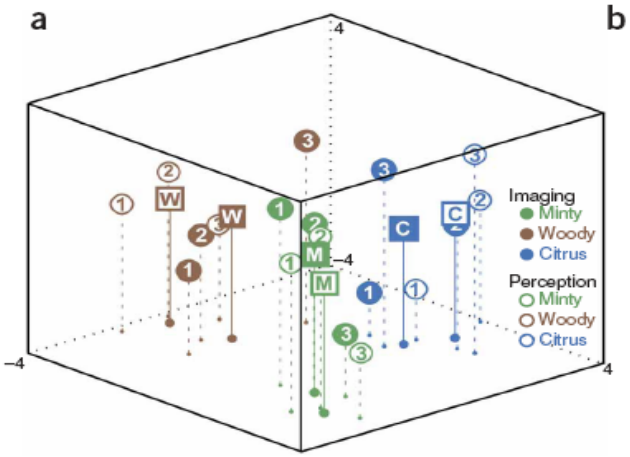


Figure: Good correlation between fMRI spatial patterns and perceived odor quality. The group-averaged fMRI data and the perceptual data were each projected onto a common three-dimensional space. Filled circles represent distributed ensemble activations. Empty circles represent odor experiences. Distances represent degrees of similarity. Squares labeled 'M' (minty), 'W' (woody) and 'C' (citrus) represent centroids of each category for the imaging (solid squares) and perceptual (empty squares) data. Degree of similarity (distance) among odor experiences is remarkably well correlated with degree of similarity (distance) among the corresponding ensemble activations. From Howard et al. 2009, with permission.

Now I can explain the significance of the fact that Howard et al. used structurally very different molecules that smell very similar (minty, woody, or citrus): that is, why they focused on a case of bad external correlation in my sense. The reason is this: because the molecular structures in each of the three families are quite different, we know that the similarities in the ensemble activations are not mere artifacts of similarities in the molecular structures of the odorants used (they were not similar); the only explanation of the correlation between ensemble activation similarity and phenomenal similarity is that ensemble activation similarity somehow plays a role in determining phenomenal similarity. As Margot (2009, 814) puts it in an overview of Howard et al.'s important new experiments:

> [Because structurally diverse chemicals were involved] the fMRI effects were not merely reflecting odorant-specific differences [and similarities]. . . . The fMRI effects unequivocally demonstrated that only the PPC ensemble activities are predictive of the category (woody, minty or citrus) of the odor that the subjects smelled. Because the chemical structure of the odors in each odor category are very different, this is strong support for the idea that the PPC codes [i. e. determines] odor quality rather than structural and chemical similarity [in the odorants tracked].[18]

The Howard study on humans is not the only study showing that similarity and difference in olfactory experience to be well-correlated with similarity and difference in ensemble activations. It is just the only one done on humans. There

---

[18] To see more clearly the point Howard and Margot are trying to make, consider an analogy. Forget about the physical basis of experience and turn to the physical basis belief for a moment. Suppose we discover that our "language of thought" expressions for dogs and cats are neurally similar: the "content-vehicles" are similar. Does this support an "internalist hypothesis" about belief-determination: does it show that these beliefs represent similar animals *because* the content-vehicles are similar? No – the orthodox view on these matters is that the intrinsic features of the neural content vehicles don't play an explanatory role in determining what they represent. The relationship is largely arbitrary. Neural similarity is incidental to content similarity (Fodor and Lepore 1999). (Compare: the public language words 'dog' and 'cat' don't represent similar animals by virtue of any similarity between them.) There is an *alternative*, "externalist explanation" of fact that our dog concept and our cat concept are physically similar. As Prinz 2006, 96 puts it (while making a different point):

> suppose concepts are learned, and suppose learning works by strengthening connections between neurons that are caused by our encounters with category instances. We see a bunch of dogs, and connections between neurons are shaped by those encounters. Then we see a bunch of cats, and the same thing happens. Chances are the resulting patterns of activation will be similar.

Now let us turn back to the case of smell experience, and to the point that Howard et al and Margot are making. We have similar smell experiences, and they are realized by similar neural patterns. Why? Here the "externalist explanation" (one that Tye and Dretske might hanker for) is unavailable. We cannot explain the match between phenomenal similarity and neural similarity by noting that the similar smell experiences are caused by similar things in the world – because in the Howard experiment they are *not* caused by similar things (odorants) in the world: see Figure 5 again. Here the only reasonable explanation of the data is some kind of broadly "internalist" one: neural similarity is not incidental to phenomenal similarity – rather it grounds phenomenal similarity. Hence the observed neat correlations.

are also many studies on animals that show this as well. Thus Youngentob and coworkers did a similar study on rats. They used the following odorants:

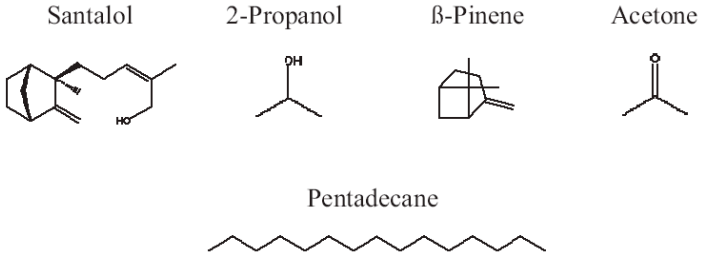


Figure 1. Chemical structures of the five odorants studied.

Now of course Youngentob and co-workers could not ask rats about the degrees of similarities among their smell experiences of these chemicals. Instead what they did was to use a confusion matrix task. The basic idea is that degree of similarity corresponds to probability of confusion. Then, using a 2-DG functional mapping technique and multidimensional scaling, they looked at the degrees of similarities among the ensemble activations set up by the odorants in the rats' olfactory bulb. The following Figure shows the remarkable correlation they found between phenomenal similarity and neural similarity:

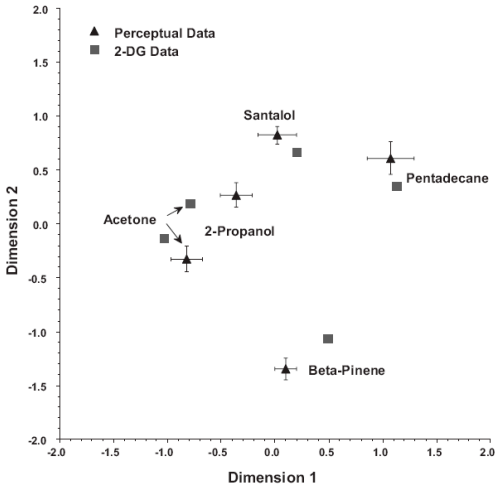


Figure 3. Alignment of the combined functional activity and behavioral odorant spaces. Coordinates are arbitrary and their scale is standardized to a common variance for the five behavioral and functional data points. The two-dimensional standard errors for the five behaviorally determined data points are shown.

Summarizing these results, Youngentob and co-workers write:

> We found a remarkable predictive relationship between the odorant-specific glomerular activity patterns and the perceptual relationship among the odorants. When the activity pattern for two odorants mapped relatively close to each other in the functional MDS space, then so did the perceptual data. Alternatively, when the 2-DG activity patterns mapped relatively distant from each other in the MDS space, then so did the behaviorally derived perceptual data. . . Our results support a combinatorial coding model in which the total pattern of bulbar activity is relevant to the production of an odorant's perceptual quality. . . . Indeed, our results show neural and perceptual relation-

> ships that could not be presumed from any prior notion of molecular similarity among the odorants. There was a greater perceptual and [neural] pattern similarity between pentadecane and santalol, than between either of these odorants and β-pinene, yet both santalol and β-pinene are bridged polycyclic compounds. . .

In my terms, what they are saying is that they found good internal correlation even when there was bad external correlation. So the results of their experiment on rats are similar to those of the experiment conducted by Howard and coworkers on humans.

3.3 Pain

Now let us leave the chemical senses, taste and smell, and turn *pain*. Suppose you have a variety of pains of different intensities in different apparent bodily locations: throbbing pains, prickling pain, stabbing pain, heat-induced pains of various intensity, and so on. On tracking intentionalism, felt pains reduce to types of bodily disturbance, just as colors reduce to reflectance properties; and every aspect of the pains reduces to some physical feature of the bodily disturbance represented by the pains. So felt location is just represented location; differences in quality among pains (prickling, stabbing, throbbing, etc.) are constituted by the quite different types of bodily disturbance they represent. Now, besides quality and location, you can also focus on the *intensity* of a pain. What (possibly complex) aspect or feature of pain constitutes its intensity? As far as I know, tracking intentionalists have not addressed this issue. But they would presumably hold that the intensity of a pain is fully determined by the intensity and size of bodily disturbance optimally tracked and hence represented.

But, in fact, even under optimal conditions, pain intensity is very poorly correlated with these factors. For one thing, psychophysics has shown that there is *response-expansion*. Even under optimal conditions, the relationship between pain intensity and bodily disturbance is described by a power function with an exponent greater than 1, where the size of the exponent differs for different kinds of stimuli. For instance, as Stevens et. al have shown (1958), doubling electric shock far more than doubles pain intensity (Figure 7a). And heat-induced pain intensity is a power function of stimulus-temperature, with a power-function greater than 1: see the figure below.

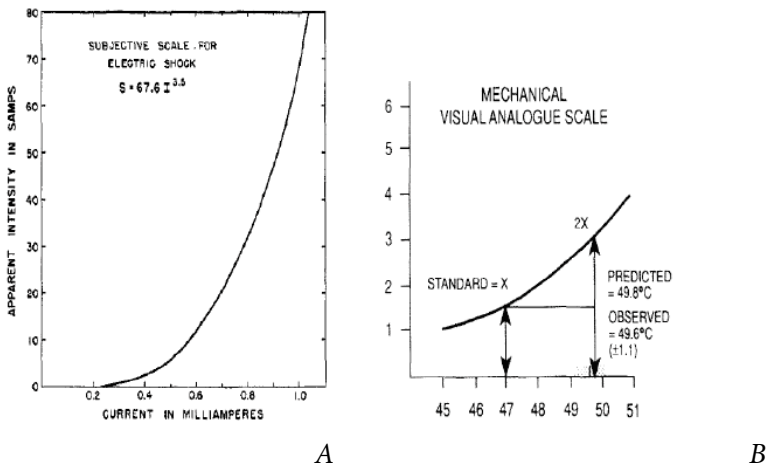


Figure 8: Psychophysical response functions for two types of noxious stimuli: electric shock (A) and nociceptive temperature (B). From Stevens 1958 and Price 1999.

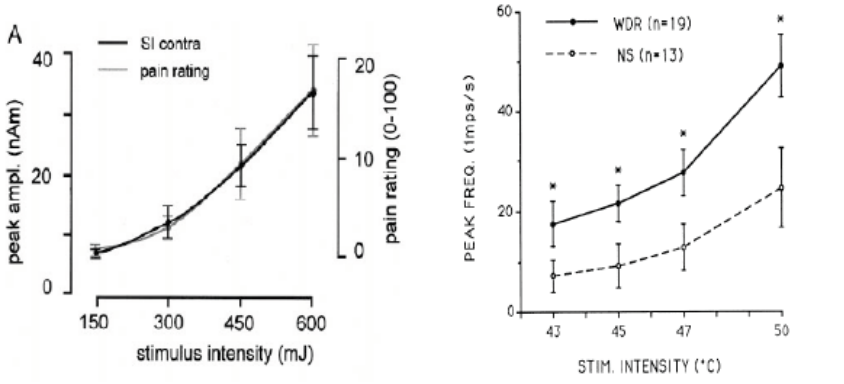
In fact, the situation is *much* more messy than these figures would suggest. Pain intensity is not only a function of stimulus intensity; it is also a function of stimulus area, in a way that cannot be codified (Price 1999).

So psychophysics has revealed that, even under optimal condition, there is a messy relationship between pain intensity and the bodily disturbance tracked. By contrast, many researchers have report perfect correlations between firing rates of neurons in the areas of the brain involved in pain and human subjects' pain ratings. For instance, using nociceptive temperatures and measuring neural activity with fMRI, Coghill et al. (1999) found linear relationships, with different regression coefficients for different areas. They write:

> Many cortical areas exhibit significant, graded changes in activation *linearly related to* pain intensity. . . . Normalized CBF differences . . . .confirm that the regression coefficient accurately describes the quantitative relationship between brain activation and perceived pain intensity. For example, the regression coefficient of the medial thalamus was 0.5, and the average psychophysical rating of 50°C was 15.6. Accordingly, the predicted activation difference between scans of 50°C stimulation and rest would approximate 0.5 * 15.6 or 7.8 (in units of normalized CBF). The observed activation difference was 7.07 (in units of normalized CBF).

Using lasers as their pain stimulus and a different technology – namely MEG – to determine neural response, Timmerman et al. (2001) found a particularly close relationship between human subjects' pain intensity and firing rates of neurons in the primary somatosensory cortex (S1) –for short *S1 firing rates* (figure 8a). As they put it, "amplitudes of contralateral S1 activity . . . match precisely the subjects' pain ratings". Kenshalo et. al (2000) also found in a single-unit study that the relationship between temperature and wide-dynanic range (WDR) neurons in monkey S1 (Figure 8b) very closely resembles the psychophysically-derived relationship between temperature and pain intensity in humans (Figure 7b). In-

deed, just as pain intensity doubles between 47C and 50C (cf. Figure 7b), so WDR response in monkeys roughly doubles between these temperatures. I should mention that some (Porro 1998, Price 1988) hold that some what determines pain intensity is firing rate and number of neurons active (the sum total or neural activity) but this issue will not matter in what follows.



A *B*

Examples of good internal correlation in the case of pain intensity. *A:* a precise match between S1 neural response and pain ratings (from Timmerman 1999). *B:* responses of monkey S1 WDR neurons to stimulus intensity.

These experiments and a slew of other empirical considerations (Price 2002) suggest that S1 plays a special role in determining pain intensity. This is the leading view I will generally assume it is correct in what follows, but none of my arguments hang on it: they would go through on a more distributed view of the neural basis of pain intensity as well (e. g. Coghill 1999).

So, in the case of pain as in the case of taste and smell, while there is bad external correlation, there is very good internal correlation. There is messy, uncodifiable relationship between the size and intensity of the various types of bodily disturbances and the S1 firing rates they set up; and in turn there is a linear correlation between these S1 firing rates and pain intensity.

The fact that when it comes to phenomenal character there is "bad external correlation" but "good internal correlation" across the various modalities makes one suspect that there is something very wrong with the radically externalist approach promoted by tracking intentionalists, according to which phenomenal character is fully determined by the external physical properties tracked by our experiences (bodily disturbance types, chemical properties, etc.) and which accords no serious role to internal factors.

In the following two sections, I will turn this suspicion into two independent empirically-based arguments against tracking intentionalism. Then I will add a third, more apriori argument against tracking intentionalism.

4 First Argument: The Internal-Dependence Argument

With the empirical background out of the way, we are now finally read to turn to my arguments against tracking intentionalism.

I call my first argument the *internal-dependence argument*. Bad external correlation together with good internal correlation suggest that tracking intentionalism

is mistaken. Phenomenal character is not fully determined by the external physical property optimally tracked. It is also shaped by postreceptoral neural processing in a way incompatible with tracking intentionalism. Tracking intentionalism takes externalism in the philosophy of mind too far.

In my first argument I will make this incompatibility precise by describing counterexamples to tracking intentionalism. The counterexamples take the form of what I call coincidental variation cases, which have the displayed structure:

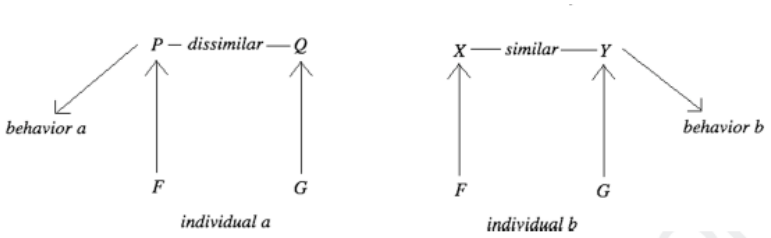


Figure 4. Coincidental variation.

In such cases, two individuals (perhaps belonging to different species) optimally track *exactly the same* external properties (*F*, *G*, . . .), but they do so *via* different neural states (*P*, *Q*, *X*, *Y*, . . .), which stand in different similarity-difference relations, and lead to systematically different sorting and other behaviors. The neural states differ in the kind of neural patterns which we saw in the previous section to be well correlated with phenomenal parameters. Analogy: a mercury thermometer and a thermoelectric thermometer have different internal states, but they optimally track the same temperatures.

You can probably guess how the argument will work. Given the empirical research covered in the previous section, and given that the individuals involved differ systematically in their neural states and behavior, it's very reasonable to think that in *some* coincidental variation cases the relevant individuals would have phenomenally different experiences. But tracking intentionalism has the completely incredible consequence that in *all* such cases the individuals involved would have phenomenally identical experiences, despite the radically neural and behavior differences, because they would sensorily represent the very same external properties. It goes totally against all of the research described in the previous section. Indeed it goes against decades of research in psychophysics and neuroscience, which show that some aspects of phenomenology can only be explained by what happens in the brain.

Now unfortunately I do not know of any clear *actual* coincidental variation cases. (If anyone does, I would like to hear about them!) For instance, as we saw before (§2), if John sees pure blue and Jane sees green-blue, the tracking intentionalist might just say that they optimally track (and hence represent) overlapping but distinct ranges of reflectances. In other words, he might say that there is not a total coincidence in what is tracked and hence represented. So there is a subtle representational difference as well as a phenomenal difference: no counterexample to tracking intentionalism. Or again: a human and a pigeon arguably have different color experiences, but they optimally track different ranges of reflectances too.

So, the cases I will discuss are hypothetical, although they are based off of actual cases. I will stipulate that in these cases there is complete coincidence in what

is optimally tracked, not just overlap. So, if we say that in such cases the individuals involved (would) have different experiences, we have clear counterexamples to tracking intentionalism. In other words, these cases succeed as counterexamples to tracking intentionalism, where actual cases seem to fall short. After I discuss some such cases, I will lay out my argument, and also say why it is superior to arguments already in the literature, like spectrum inversion and Levine's argument about the fish, which I criticized already (§2).

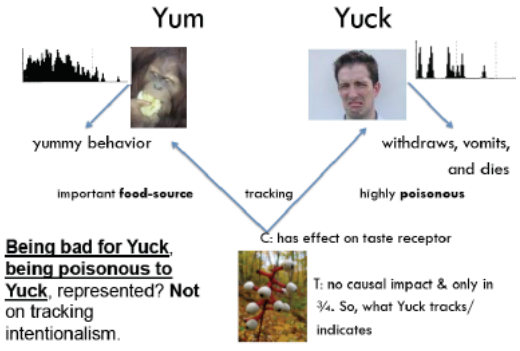
I will save objections objection until §7.

4.1 First Apparent Counterexample: Yuck and Yum

My first case can be introduced *via* an actual case. The berries *actaea pachypoda* (Doll's-eyes) is highly poisonous (and bitter in taste) to humans, but harmless to birds, the plant's primary seed dispensers. They eat it up without problem. Now it is reasonable to think that while the berries taste horribly bitter to us, they taste different to the birds.

Is this an actual counterexample to tracking intentionalism? No – unfortunately for me humans and the birds differ at the receptoral level too, so that the brain states that realize their experiences of the berries are caused by different ranges of chemical properties. So the tracking intentionalist can say that the phenomenal difference is matched by a representational difference. Analogy: humans and pigeons have different color experiences, but they also optimally track and hence represent different ranges of reflectances. So here we have no counterexample to tracking intentionalism.

But with a small twist we do get a counterexample to tracking intentionalism. Just consider a hypothetical *coincidental variation case* in which the brain states of the two individuals involved do optimally track the very same chemical property of the berries. It's still reasonable to think that berries taste differently to them, but tracking intentionalism is inconsistent with this verdict.



In a bit more detail: suppose Yuck and Yum belong to different species that evolved in separate environments containing the berries mentioned above: Doll's eyes. Now you might suppose that Yuck is an actual human – me or you – and Yum is some hypothetical creature. Or you might suppose that Yuck and Yum both belong to hypothetical, human-like species. It does not matter. In any case, the berries are extremely poisonous to Yuck. But in Yum's environment, the berries are a very important foodsource, since other foodsources are scarce. So Yum's species evolved immunity to the berries. In addition, when Yuck and Yum taste the berries, their taste systems undergo radically different ensemble activa-

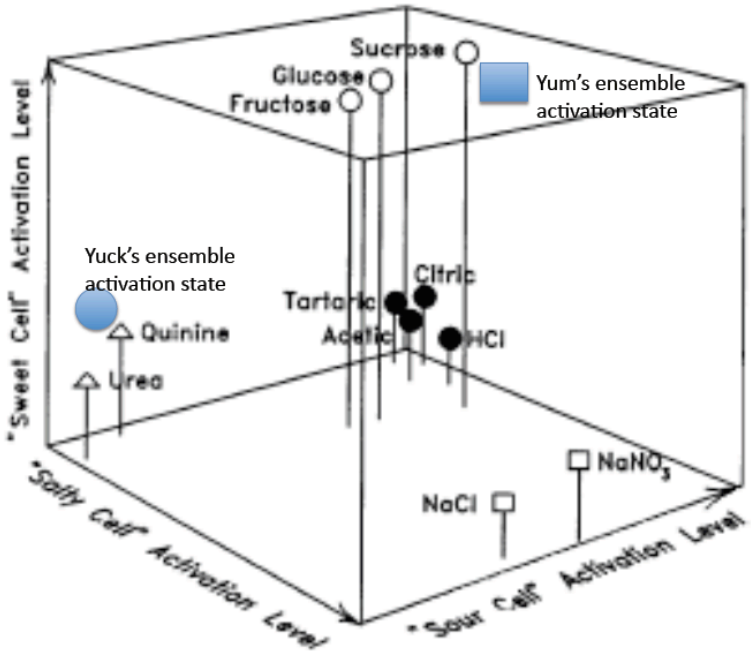
tions (across-fiber patterns). They also innately disposed respond to their tastes experiences with radically different behaviors. For instance, Yuck vomits and withdraws from it violently, while Yum is drawn to it, rubs his tummy, and so on.

Despite these differences, we can stipulate that Yuck and Yum are similar at the receptoral level, so that, when they taste the berries, the postreceptoral ensemble activations in their taste systems, although different, optimally track the very same chemical property of the berries, *C*. (I do not say that the fact that their receptoral states carry the same information about the environment shows that their postreceptoral states further downstream do; I am stipulating that the postreceptoral states carry the same information about the environment, at least on the tracking intentionalist's psychosemantics.) Likewise, the ensemble activations in their smell systems optimally track the same chemical property of the odor released by the berries.

Now it might be said that Yuck but not Yuck tracks the poisonousness-to-Yuck of the berries. We can rule that out by supposing the poisonousness of the berries is grounded in chemical property, *T*, different from *C*, which is undetectable by the taste system and which only three-fourths of the berries possess.

I said that Yuck and Yum undergo different ensemble activations in response to the berries. Let me be more specific. Suppose that the notorious poison dart frog is highly poisonous to both Yuck and Yum. Suppose further that, when Yuck tastes berries, the ensemble activation he undergoes is quite similar to the one he undergoes when he tastes the dart frog. By contrast, when Yum tastes the berries, the ensemble activation Yum undergoes is radically different from the one is undergoes when he tastes the dart frog, and much more like the one he undergoes when he tastes yummy bananas.

In general, the set-up is that the ensemble activation that the berries produce in Yuck is similar to those which he undergoes when he tastes things that presumably taste bad or bitter to him, whereas the ensemble activation that the berries produce in Yum is similar to those which he undergoes when he taste things that presumably taste good (e. g. sweet) to him. Diagrammatically:



That, then, is the case described in non-phenomenal terms. The crucial question for the tracking intentionalist is whether Yuck and Yum would have different taste experiences or the same taste experiences of the berries.

I think we should say that they would have different experience experiences. We have seen that resemblances and differences in taste quality is much better correlated with resemblances and differences in ensemble activations than with resemblances and differences in the chemical properties optimally tracked. So the ensemble activation differences between Yuck and Yum *are very good evidence* that they have phenomenally different taste experiences of the doll's eyes berries. By contrast, the fact that they optimally track the same chemical property of the berries is very *poor* evidence of phenomenal sameness.

More specifically: when Yuck tastes the berries his ensemble activation state is very close to those which he undergoes when he tastes the poison dart frog and other characteristically bad-tasting things, whereas when Yum tastes the berries the ensemble activation he undergoes is similar to those which he undergoes when he tastes bananas and other characteristically sweet things (see again the figure above). Given that similarities and differences among ensemble activations are the only things in the physical world that correlate well with similarities and differences among taste experiences, Yuck's taste experience of the berries is probably similar to his experience of the poison dart frog and other characteristically bad-tasting things, whereas Yum's taste experience of the berries is probably similar to his taste experience of bananas and other characteristically good-tasting things.

The behavioral differences between Yuck and Yum suggest an independent argument for the same verdict. When Yuck has the berries, he exhibits certain innate responses: he vomits, withdraws, and treats them like other things that presumably taste bad (e. g. bitter) to him. When Yum has the berries, he seeks more, and treats them like other things that presumably taste good (e. g. sweet) to him, like bananas. Even if we're not behaviorists or functionalists, this suggests different experiences of the berries, just as humans and birds to in the actual world.

Note well that I only say that, given what we know about the physical basis of taste experience, it is reasonable to suppose that there are phenomenal differences between Yuck and Yum. This is all that is required by the argument. I do not say that we can read off from the physical facts exactly what their taste experiences are like. Nor do I say that we can read this off more specifically just from the character of their neural states.

So the only reasonable verdict is that Yuck and Yum would have different taste experiences of the berries. But tracking intentionalism delivers the incredible verdict that Yuck and Yum would have phenomenally identical taste experiences of the berries, despite the vast neural and behavioral differences between them when they taste the berries.

To see this, notice first that, by stipulation, Yuck and Yum are in neural states which, although different, track exactly the same external chemical property, C, of the berries. Further, *optimal conditions* obtain. It is not as if one is a genetic freak, or has a malfunctioning taste system. On the contrary, their taste systems, although different, are both working as they were designed by evolution to work. Further, their innate behavioral dispositions in response to the berries, although

different, are adaptive, given the difference in the biological significance of the berries to them. On tracking intentionalism, the symmetry in the optimal tracking facts between Yuck and Yum means that their taste experiences of the berries are representationally exactly the same; in particular, they both sensorily represent the very same chemical property, C, of the berries. In general, their inner ensemble activations, although different, represent the same chemical properties of foodstuffs.

Consider some analogies: despite their different internal workings, a thermoelectric thermometer and a mercury thermometer are representationally the same. Or again: English words stand in certain similarity and difference relations, but this is incidental to what they represent. Thus, 'river' and 'liver' are similar but represent quite different things. In a different language, the words for rivers and livers might instead be radically different from each other; but they would be representationally the same, representing rivers and livers. Likewise, on tracking intentionalism, even though Yuck and Yum's ensemble activation states stand in quite different relations of similarity and difference, they represent the very same external chemical properties: on this view the similarity/difference structure of those states in incidental to what they represent. Further, on tracking intentionalism, the behavioral differences between Yuck and Yum do not show that their inner states differ in what they represent. This is a natural idea and we will later see if it might help intentionalists with this sort of case (see the discussion of Lewis, Jackson and White in the conclusion). But *tracking* intentionalists cannot accept it: they do not advert to behavioral dispositions in a serious way in their account of what determines representational content (except Tye says that they have to be appropriate in order that optimal conditions obtain).

Now I do not say that the verdict of tracking intentionalist that Yuck and Yum's ensemble activation states are representationally identical is *by itself* bad. Ensemble activation states might represent external chemical properties *in some sense*, even if there is "bad external correlation", that is, even if the resemblances and differences among them are not matched by resemblances and differences among the chemical properties. On many views, anything can represent anything: the connection between the intrinsic character of content-vehicles and what they represent is arbitrary.

But recall that on tracking intentionalism phenomenal character is also fully determined by representational content. Once we add this to the tracking intentionalist's commitment of total representational sameness between Yuck and Yum, we do get a bad verdict: that Yuck and Yum have phenomenally identical experiences of the berries, despite the vast neural and behavioral differences between them. *This* verdict is radically implausible. Again, given the neural and behavioral arguments I gave above, by far the most reasonable verdict is that they would have phenomenally different experience of the berries (yuck and yum), just as humans and birds do in the actual world.

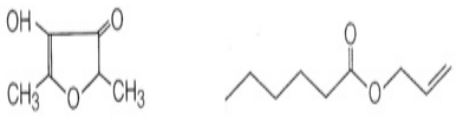
Let me conclude by saying that this apparent counterexample to tracking intentionalism is not far fetched at all, in contrast to spectrum inversion, inverted earth, and the like. The case is a twist on an actual case: the berries doll's eyes is actually poisonous and bitter-tasting to humans but an important food-source and presumably good-tasting to other animals.

4.2 Second Apparent Counterexample: Harold and Henry

My second apparent counterexample concerns an actual case that we discussed earlier. It can be represented like so:

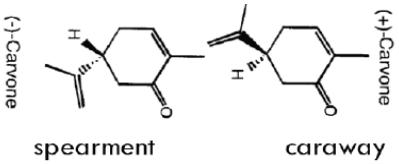
Henry P ------different--------X C* -------similar -------C

Harold P -------similar ------------P* S ------different--------C



X

Y



The first and second molecules (I'll call them X and Y; Y is allyl hexanoate but I do not know what X is called) are clearly more different from each other than the second two (namely -carvone and +carvone). Indeed, the second two molecules are enantiomers of carvone: stereoisomers that are *mirror images*. Now suppose an ordinary human in the actual world – call him Harold – smells the molecules. As a matter of empirical fact, the first two produce very similar (but non-identical) pineapple smells in him, whereas the second two smell like spearmint and caraway, respectively. In other words, this is a striking case of "bad external correlation".

Now as we saw Linster found that stereoisomers that smell different to humans also produce different ensemble activations in the olfactory bulb of rats. We can conjecture that -carvone and +carvone produce very different ensemble activations in Harold's olfactory system, namely S and C. We also saw previously that Howard and Youngentob showed that in general similarities between smells are remarkably well correlated in humans and animals with similarities between ensemble activation states, even in cases where the chemical properties they track are radically different. Harold has similar pineapple experiences of X and Y. So we can also reasonably conjecture that X and Y, despite being chemically very different, produce very similar ensemble activation states in Harold, call them P and P*.

Now consider a nearby counterfactual situation in which we evolved a slightly different olfactory system. It does not matter why: matter the difference is an adaptation to the environment, as in the case of Yuck and Yum. Or maybe the difference is due to fluke differences in the evolutionary process. In any case, consider a particular individual, *Henry*, in that situation. Whereas in the actual world the mirror-image enantiomers -carvone and +carvone optimally cause in Harold quite different ensemble activations S and C, in the counterfactual situation they optimally case in Henry very similar ensemble activations C and C*. Further, whereas in the actual world X and Y optimally cause in Harold very similar ensemble activations P and P*, in the counterfactual situation they optimally cause in Henry very different ensemble activations P and X.

In addition, there are behavioral differences between Harold and Henry. Harold sorts x-carvone and y-carvone apart and has an easy time distinguishing them. Henry can distinguish them but has a hard time doing so; and he sorts them together. Harold sorts X and Y together; Henry sorts them apart.

Now in the actual world Harold's smell experiences of the first two molecules are more similar than his smell experiences of the second two. We don't need an argument to show this: it is an actual fact. What about Henry's experiences? I think we should say that his smell experiences of these molecules are (that is, would be) phenomenally different from Harold's. There are two arguments for this, one neural and the other behavioral.

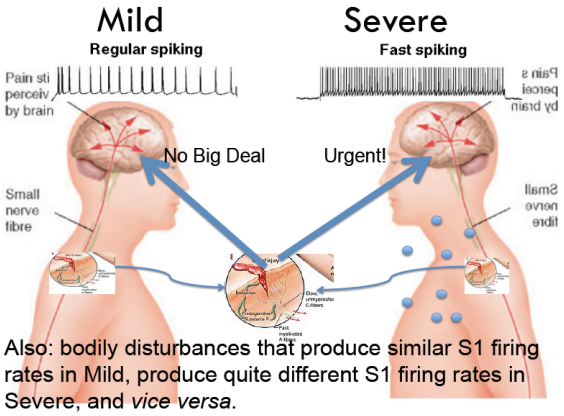
First, what explains why *Harold's* smell experiences of the first two molecules are more similar than his smell experiences of the second two? Not that the first two molecules are more similar than the second two (they aren't)! The only explanation is bring in internal, postreceptoral factors. The explanation is that the ensemble activations produced in Harold by the first two are more similar to those produced in him by the second two. At the very least, the resemblance pattern of the internal ensemble activations *has something to do with it*. This is virtually shown by the Howard and Youngentob experiments. Now the ensemble activations produced in *Henry* by first two molecules are *not* more similar to those produced in him by the second two. Rather, they are vastly more different than those produced by the second two. So it stands to reason that, by contrast to Harold, Henry's smell experiences of the first two are *not* more similar than his smell experiences of the second two. Maybe his smell of the first is pineapple-like, his smell experience of the second is musk-like, and his smell experience of the second two are both caraway-like. It does not matter. In any case, there are phenomenal differences between Harold and Henry's smell experiences. A second argument for this verdict is based on the hardwired, innate behavioral differences between Harold and Henry with respect to the molecules.

So the claim that Henry would have different smell experience of the molecules than Harold is firmly rooted in empirical research. It looks as if ensemble activation patterns are the only thing in the natural realm that might explain smell quality. As Margot put it, "[ensemble activation] codes [i. e. determines] odor quality rather than structural and chemical similarity [in the odorants tracked]". And Henry differs greatly from Harold in his ensemble activations, as well as in his behavior.

But tracking intentionalism delivers the incredible verdict that there are no phenomenal differences between Harold and Henry's smell experiences. On this view, the neural and behavioral changes do not make for any representational differences between Harold and Henry. The switches in the neural content-carriers, with their very different intrinsic characters, makes no differences to the externally-determined contents-carried. So, in spite of all the evidence to the contrary, they do not make for any phenomenal differences, according to tracking intentionalism. For, as in the case of Yuck and Yum, we can stipulate that Harold and Henry's ensemble activation states, although different, track under optimal conditions the very same (ranges of) chemical properties. The verdict of tracking intentionalism that Harold and Henry have exactly the same smell experiences, despite the vast neural and behavioral differences between them, flies in the face of the empirical research on the physical basis of smell quality, which locates that basis in the brain.

4.3 Third Apparent Counterexample: Mild and Severe

Here is my final coincidental variation case, one involving pain:



Mild and Severe belong to different species that evolved separately. Each is subject to a lesion in his arm. In Mild, the lesion is not very dangerous. So under optimal conditions *it* produces a mild S1 firing rate. By contrast, species to Severe belongs evolved in an environment in which a lesion like this is extremely dangerous, because in this environment a certain sort of airborne bacterium is present in abundance. In consequence, under optimal conditions *it* produces an S1 firing rate in which is, say, twice greater that produced in Mild. As a result, Severe attends to his leg with greater urgency than does Mild. Further, suppose that bodily disturbances that produce similar S1 firing rates in Mild, produce quite different S1 firing rates in Severe, and *vice versa*. For instance, suppose both Mild and Severe sustain moderate blunt trauma. We might suppose that the S1 firing rate produced in Mild by the lesion is similar to that produced by the blunt trauma, whereas the S1 firing rate produce in Severe by the lesion is much greater than that produced by the blunt trauma. Consequently, Mild treats the lesion and the blunt trauma with similar urgency, whereas Severe treats with lesion with much greater urgency than he treats the blunt trauma.

We have seen that in humans there is linear correlation between pain intensity and S1 firing rates, whereas there is no codifiable relationship between pain intensity and the bodily disturbance optimally tracked. So *the differences between Mild and Severe in firing rates of primary somatosensory neurons is very good evidence that they have pains of different intensities, and the fact that they optimally track the same bodily disturbance is very poor evidence that they have the same pain experience*. The behavioral differences between mild and severe suggest the same verdict.

Now it might be said that this case is weaker than the other cases. For can we really assume that absolute differences in S1 firing rates *across species* go with absolute differences in pain intensity? Maybe different S1 firing rates can realize the pain of the same intensity across species: the same pain intensity is multiply realizable. Maybe that is what is going on in the case of Mild and Severe.

I have a few points in reply. (i) I'm *not* assuming that absolute differences in S1 firing rates *across species* go with absolute differences in pain intensity. I'm appealing to relative differences in S1 firing rates, and I'm also appealing to behavioral differences. We know pain intensity is only well correlated with S1 firing rates. The S1 firing rate produced in Mild by the lesion is similar to that produced by the blunt trauma, whereas the S1 firing rate produce in Severe by the lesion is much greater than that produced by the blunt trauma. Consequently, Mild treats the lesion and the blunt trauma with similar urgency, whereas Severe treats with

lesion with much greater urgency than he treats the blunt trauma. So it's reasonable to think that Mild's experiences of the lesion and the blunt trauma are phenomenally similar as regards intensity, whereas Severe's experiences of the same stimuli are phenomenally different as regards intensity. This means that there are phenomenal differences between Mild and Severe's experiences. This is enough to refute tracking intentionalism, we shall see. In short, it's not at all reasonable to treat the case of Mild and Severe as a case of alternative realizations of the very same experiences: they have different experiences. It is not like a stock case of multiple realizability at all. In the stock cases of multiple realization – for instance a human and Martian – there is supposed to be behavioral *similarity* and only superficial internal differences, whereas in the case of Mild and Severe there are profound neural and behavioral differences, making the verdict of different experiences the only reasonable one. (ii) We can assume that Mild and Severe belong to a common lineage that branched, so that they are very similar, aside from the differences I have mentioned. Then the neural and behavioral differences between them are even stronger evidence of phenomenal differences. (iii) Finally – although this is not required by my argument as we have just seen – I do not think it is crazy to assume that absolute differences in S1 firing rates *across species* go with absolute differences in pain intensity. Something in the physical world has to determine absolute pain intensity. Something also has to fix truths like "Severe's experience is twice greater in intensity than Mild's" or "Mild's two pains are more similar than Severe's two pains as regards intensity". These could be true, even if we cannot know *for certain* that they are true. Something in the physical world fix these truths. This is so on dualism as well as physicalism. (On Dualism the physical facts fix the phenomenal truths via contingent psychophysical laws.) What in the physical world could fix these truths? One idea is: facts about functional role, principally facts about behavior. (Under dualism, the idea would be that there are immensely complicated psychophysical laws going from such facts to facts about absolute pain intensity.) But is it really plausible to think that differences in absolute pain intensity are fixed by "degrees of urgency" in behavior (whatever that might mean) – so that there is one level of urgency that goes with this pain intensity, another that goes with that pain intensity, and so on, across all species? And what could be the functional-behavioral ground of one creature's pain being *twice more intense* than another's?[19] The only other alternative is that these truths are fixed by neural facts; and it would appear that the most likely candidates are facts about S1 firing rates (or the like).

In any case, whatever one's exact views about the ground of phenomenal character, the vast neural and behavioral differences between Mild and Severe, together with what we know about the physical correlates of pain intensity, mean that the only reasonable verdict about this case is that they have different pain experiences. But tracking intentionalism delivers the opposite verdict. For Mild and

[19] Shoemaker has long held that facts about *resemblances a*mong color experiences are grounded, not in resemblances among neural states, but in functional facts about what beliefs they're apt to cause and ultimately in facts about behavior. But he nowhere offers a functionalist account of ratio scaling facts- which seems to me pretty impossible. Also the account is problematic. (i) The account has problems with animals or babies who don't have beliefs about resemblances among their experiences. (ii) The account faces serious problems with across-individual comparisons (like that between Mild and Severe): for instance, their experiences don't cause anyone to judge one pain is twice greater than another, etc. See Levine (1988) for criticisms of Shoemaker's attempt to solve the problem with across-individual comparisons.

Severe's S1 states, although different, optimally track the very same bodily disturbances. In generally, there is symmetry between Mild and Severe as regards the optimal tracking facts. So the representational contents of their pain experiences are the same. I have no objection to the claim that in some sense Mild and Severe's pains represent bodily disturbances. As we saw before, there is a bad correlation between S1 firing rates and the size and intensity of the bodily disturbances tracked; but this does not mean that S1 firing rates cannot represent the size and intensity of the bodily disturbance tracked. On most theories of representation, anything can represent anything. But, on tracking intentionalism, the representational contents of Mild and Severe's pain experiences are completely the same, and this has a further consequence that I do object to: namely, that they are completely phenomenally the same. This verdict is totally implausible for the reasons I have given.

4.4 The official internal-dependence argument

Of course, such cases could be multiplied indefinitely. To refute tracking intentionalism, only one counterexample is required. So the best way to state the argument is as follows.

1 If tracking intentionalism is true, then in *every* possible coincidental variation case, the right verdict is Same Experiences.
2 But it is much more reasonable to suppose, in at least *some* coincidental variation cases the right verdict is Different Experiences; call this *Internal-Dependence*.
3 So tracking intentionalism is (probably) mistaken.

Since my argument only requires that there be one coincidental variation case in which the right verdict is Different Experiences, you shouldn't get too bogged down in any particular case. If you don't think one case works as a counterexample, there is bound to be another that does.

I will consider some desperate objections to this argument later on (in §7) but basically I think that the argument is totally persuasive. All it's trying to show is that it's more reasonable to reject tracking intentionalism than to accept it, and that's because it's more reasonable to accept Different Experiences in some coincidental variation cases than to accept Same Experiences in all of them. This just seems like a no brainer in view of the empirical evidence we looked at earlier. To suppose that Same Experiences is the right verdict in all coincidental variation cases is to suppose that the external physical properties optimally tracked (chemical properties, bodily disturbances) are the *only determinate of phenomenal character*. Given that phenomenal character is in general very poorly correlated with the external physical property optimally tracked ("bad external correlation"), and is much better correlated with neural states in the brain ("good internal correlation"), this would just be preposterous. Also, just imagine meeting Yuck and Yum, or Harold and Henry, or Mild and Severe. Even a tracking intentionalist would have to admit that they have different experiences, given all the evidence. Can you imagine him denying this in the name of a speculative philosophical theory? So he ought to give up his view.

Let me make a few clarifications about premise 2, Internal-Dependence. This is just the claim that in some (i. e. at least one) coincidental variation cases the

right verdict is Different Experiences. This is very weak: it is weaker total internalism, the thesis that internal factors and internal factors alone totally determine phenomenal character. Yet Internal-Dependence is enough to refute tracking intentionalism. Internal-dependence also does not bring with it commitment to any particular theory of experience. In particular, it does not require type-type identity theory of the kind defended by Block, McLaughlin, the Churchlands, and so on, even if it might naturally suggest such a theory. It is totally consistent with functionalism and sensorimotor views. For the individuals in my coincidental variation cases are stipulated to differ in functional organization and in action-oriented respects, as well as differing in neural processing. So one might accept Internal-Dependence, and say that the neural matters to the phenomenal, but only because it matters to functional role or sensorimotor profiles, which is what really constitutes phenomenology. Finally, and relatedly, internal-dependence does not require that there be any easily-stated "psychophysical linking propositions" (decoding algorithms) going from neural parameters alone to phenomenal parameters

It should be obvious how the internal-dependence argument is better than the usual arguments against tracking intentionalism. Some (Horgan, perhaps Block) argue tracking against intentionalism on the basis of the "intuition" that intrinsic duplicates cannot differ phenomenally. I have already said briefly why I think this argument fails. I think we do not have apriori justification for believing such a thing. Of course, my internal-dependence argument avoids this problem because it doesn't require on our having such apriori justification. Others argue against tracking intentionalism on the basis of spectrum inversion across individuals (or, as in the case of Levine's fish, spectrum inversion within a given individual) who optimally track the same external properties. But this argument just relies on the untutored intuition that such a case is possible. The tracking intentionalist can just reject this intuition. Indeed, physicalists of all people should not put any stock in such intuitions, for Zombies and spectrum inversion among total physical duplicates are intuitively possible too. My internal-dependence argument did not rely on any such intuitions; indeed it did not rely on intuitions at all. Rather, first I described, in purely physical terms, some situations that are definitely possible, namely coincidental variation cases. Then I argued on empirical grounds that in some such situations the relevant individuals would have different experiences, contrary to tracking intentionalism.

I have focused on Dretske and Tye. In fact, in §2 I provisionally defined tracking intentionalism as the disjunction of their views. But the internal-dependence argument potentially undermines a whole range of externalist-reductive versions of intentionalism. Indeed, we might adopt a broader definition of tracking intentionalism. We might say that *tracking intentionalism* is any version of reductive intentionalism that delivers the verdict of Same Experiences in coincidental variation cases because in *some* sense the individual involved "track" the same properties. I think that many (including Byrne and Hilbert and Hill) are committed to tracking intentionalism in this broad sense about at least some sense modalities even if they do not have developed theories about how to spell out the generic "tracking" relation. But tracking intentionalism must be rejected if we accept In-

ternal-Dependence: it's more reasonable to accept the verdict of Different Experiences in coincidental variation cases.[20]

Of course, there are potential objections to the Internal-Dependence Argument; but I will address objections to this argument and my other arguments against tracking intentionalism later on (in §7), after I have put them all on the table.

5 Second Argument: The Structure Argument

My internal-dependence argument against tracking intentionalism concerned hypothetical coincidental variation cases. By contrast, my *structure argument* concerns actual cases. It focuses on our actual judgments about the phenomenal structure of our experiences. The argument is really an argument-schema: if

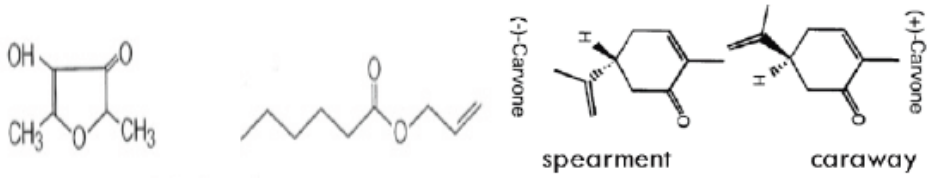
[20] Let me say why I think Byrne and Hilbert are committed to tracking intentionalism in this sense. They advocate (1997, 2003) standard intentionalism about color experience: the phenomenal character of experience is fully determined by the representation of color (together perhaps with level of illumination). And they identify colors with response-independent physical properties of surfaces. So they hold that the phenomenal character of experience is fully determined by the representation of response-independent properties of surfaces. Now they do not discuss other sense modalities in detail. But they say (Byrne and Hilbert 2003, 59) that their views of color and color experience stand or fall with similar views about other sorts of experiences and sensible properties. So they would presumably accept intentionalism about taste and smell experience, and identify the represented properties with response-independent physical properties of foods or odors (see their "Basic Sensible Properties and the Structure of Appearance"). I think this means that they must accept "tracking intentionalism" about taste and smell in the broad sense that I introduced in the text: they must accept a version of intentionalism about taste and smell that entails that Yuck and Yum and Harold and Henry have exactly the same taste or smell experiences, in which case their view is vulnerable to the Internal-Dependence Argument. For Yuck and Yum and Harold and Henry bear all the relevant naturalistic relations ("tracking relations" broadly construed) that might ground sensory representation to the same response-independent chemical properties. Now Byrne and Hilbert (2003) express skepticism about all *existing* reductive theories of representation; but, absent some argument to the contrary, *any* reasonable theory of how we might represent such properties would entail that Yuck and Yum and Harold and Henry represent the same chemical properties of foods and odors (and hence the same taste and smell properties, on a response-independent physicalist account of these). So, given their intentionalism, they are committed to the claim that they have exactly the same experiences – which I have argued is not a reasonable position. If Byrne and Hilbert's views on smell are parallel to their view on color, their view on smell is also vulnerable to the Structure Argument to be discussed in the next section. (I am not sure what they would say about *pain* – it would be interesting to know - but if their account of pain is parallel to their account of color and color experience it may be vulnerable to my arguments in the case of pain as well.)

Now I turn to Chris Hill. I think Hill is committed to tracking intentionalism, in the broad sense I introduced in the text, at least in the case of pain. In *Consciousness* he writes 'the phenomenal character of an experience is the set of qualia that the experience represents as instantiated by objects of awareness' (148). This makes him a total intentionalist about pain experience; it leaves no room for the intrinsic features of the neural content carriers (S1 states in the brain) to determine the phenomenal character of pain. He also says that in the case of pain the represented properties (the "qualia") are nothing but properties involving actual or potential bodily damage. Yuck and Yum bear the same relevant naturalistic relations to the same such properties, by stipulation. So it looks as if Hill is committed to the verdict that they represent the same such properties and hence (given his total intentionalism about pain) have phenomenally identical pain experiences. This makes him a tracking intentionalist about pain in the broad sense introduced in the text. And it makes him vulnerable to the Internal-Dependence Argument in the case of Mild and Severe: for it is not at all plausible that they have exactly the same pain experiences, given what we know about the physical basis of pain experience.

tracking intentionalism is true, then certain of our judgments about phenomenal structure are false; since they are true, tracking intentionalism is false. Perhaps the best way of stating the basic point is by saying that *there is no credible version of tracking intentionalism on which our structure judgments come out true.* I will illustrate the argument with a couple of examples. Then I will rule out some replies. Finally I will say how the argument goes beyond a well-known argument discussed due to Hardin and Byrne and Hilbert concerning the phenomenal structure of colors.

5.1 An Example Concerning Smell

My first example is one we have already encountered. Recall the following molecules:



A Z C C*

Here A, Z, C and C* are names for the chemical properties tracked by Harold's smell experiences of these molecules.

As a matter of empirical fact, Harold will report:

> [S] My smell experiences of the first two molecules resemble more than my smell experience of the second two molecules.

In other words, this is a case of bad external correlation. Yet this introspective report is surely true: the first two smell experiences are of a pineapple-like smell, whereas the second two are of spearmint and caraway.

But on tracking intentionalism the introspective report *appears to* come out false. Here is a *prima facie* argument for this claim:

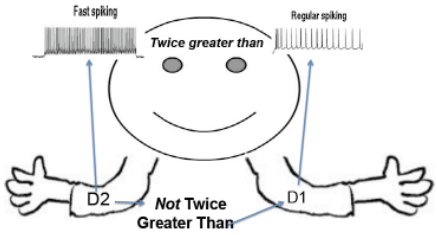
> 1 On tracking intentionalism, phenomenal facts are facts about the properties represented by experience, not facts about the neural content carriers (by definition).
> 2 So if tracking intentionalism is true, Harold first two smell experiences resemble more than his second two iff the smell properties *represented by* the first two resemble more than those *represented by* the second two. (As Byrne puts it in the case of color experience: "Why is *the experience* as of a teal object similar to *the experience* as of a turquoise object? Because *teal* is similar to *turquoise*"[21])

[21] See Byrne 2003, sect. 2. I do not suggest that on tracking intentionalism the resemblance needs to be represented by the experience. If Harold has the smell experiences sequentially there is no time at which he might represent a resemblance relation among the smell qualities because the smell

3 Optimal conditions obtain: So, on tracking intentionalism, these represented smell properties are identical with chemical properties actually instantiated by the odor clouds before Harold's nose: A, Z and C and C*.[22]
4 So if tracking intentionalism is true, Harold's introspective report is true iff A and Z resemble more than C and C*.
5 But this is evidently not true: there is no acceptable precisification of 'x and y resemble more than w and z' relative to which this predicate is satisfied by A and Z and C and C* (in that order).
6 So if tracking intentionalism is true, then Harold's introspective [S] report on the resemblances of his smell experiences is false.
7 This is absurd: surely [S] is true.
8 C. Tracking intentionalism is false.

Now I should say that this argument is not meant to complete my Structure Argument. It is just meant to provide an initial, prima facie reason for thinking that if tracking intentionalism is true then our structure judgments are false. The crucial claim is that *there is no credible version of tracking intentionalism on which our structure judgments come out true.* The tracking intentionalist might find some fault with the above argument but then he would have to provide some account of our structure judgments, consistent with his tracking intentionalism, on which those judgments come out true. I will argue (in §§5.2-5.3) that he cannot provide such an account.

5.2 An Example Concerning Pain



Suppose Max undergoes bodily disturbance D1 in his left arm and then undergoes bodily disturbance D2 in his right arm. Suppose that D2 is greater than D1 but there is no sense in which it is twice greater. But suppose that, because of what I previously called response expansion, the S1 firing rate optimally caused by D2 *is* twice greater than that optimally caused by D1. Then Max will say:

---

qualities are never presented to him all together at the same time. Yet on tracking intentionalism the resemblance order among the smell experiences derives from the resemblance order of the represented smell qualities. See Byrne 2003, sect. 6.

[22] Here A, X and C and C* are the disjunctions of the chemical properties that would cause under optimal conditions Harold's smell experiences – and hence the disjunctions that on tracking intentionalism are represented by those smell experiences. (The relevant properties are disjunctions because there is "metamerism" in the domain of smell as well as the domain of color.) The chemical properties represented in the figure are disjuncts of these disjunctions.

[P] My right arm pain is twice great than my left arm pain

He will say this even though D2 is not twice greater than D1. In other words, this is a case of what I previously called bad external correlation. Of course this case is schematic, but there are actual cases of this kind (see Figure 8 in §3).

Now consider the following initial, prima facie argument against tracking intentionalism:

1 On tracking intentionalism, phenomenal facts are facts about the properties represented by experience, not facts about the neural content carriers.
2 So if tracking intentionalism is true, the pain in Max's right arm is roughly twice greater than that in his right if, and only if, the pain property represented in his right arm is roughly twice greater than that represented in the left.
3 Optimal conditions obtain: So, on tracking intentionalism, these represented pains are identical with the actual disturbances (types), D2 and D1.
4 So if tracking intentionalism is true, Max's introspective report "the pain in the left arm is roughly twice greater than that in my right" is true iff D2 is roughly twice as great as D1. (For the claim that it makes sense to say that properties stand in ratio relations see Suppes, Mundy 1987, Swoyer 1987.)
5 But D2 is not roughly twice as great as D1.
6 So if tracking intentionalism is true, then Max's introspective report [P] on the phenomenal relation among pains is false!
7 But that's absurd. Surely, Max's judgment [P] about his own pains is true. We are very good at making such judgments (Price 1999; Myles 1999) about pains and ratio-scaling of sensory intensity is a staple of psychophysics (Gescheider 1997).
8 C. So, tracking intentionalism is false.

Now [P] involves ratio-scaling of pain intensity. As I have just said, there is strong evidence that ratio scaling of pain intensity makes sense. Still I suppose that the tracking intentionalist could bite the bullet and say that [P] is false on the grounds that ratio scaling of pain intensity makes no sense. But this would hardly get him out of the woods. For, as I have said, the structure argument is really an argument schema: if tracking intentionalism is true, then certain of our judgments about phenomenal structure are false; since they are true, tracking intentionalism is false. If you do not like the present example involving pain, we can just switch to another example, for instance the previous example involving smell.

As with my previous argument about Harold, this argument is not meant to complete my Structure Argument. It is just meant to provide an initial, prima facie reason for thinking that if tracking intentionalism is true then our structure judgments are false. In response, the tracking intentionalist might try to provide some account of our structure judgments, consistent with his tracking intentionalism, on which those judgments come out true. But I will now argue that he cannot do so.

5.3 A Response-Dependent Account Unavailable to Tracking Intentionalists

So let me address a reply on behalf of the tracking intentionalist (due to Jonathan Cohen in discussion). To illustrate the reply, consider Harold again. He comes out with introspective report [S], which says that his first two smell experiences are more similar than his second two. The reply grants that the tracking intentionalist must say that this claim is true only if the chemical properties *represented by* his first two smell experience resemble more than those *represented by* his second two: that is, just in case A and Z are more similar than C and C*. But on this reply there is a sense in which A and Z *are* more similar than C and C*, specifically, in their neural effects: the ensemble activations produced by A and Z are more similar than those produced by C and C*. In this way, by giving an account of the truth-conditions of Harold's report that brings in internal, neural states, the tracking intentionalist can say that Harold's introspective report is absolutely true.

Likewise for Max's introspective report about his pains [P]. The idea is that the tracking intentionalist should agree that it is true just in case D* is twice greater than D. But he can say that D* is twice greater than D, in the sense that the S1 state it produces involves a firing rate that is twice greater than that involved in the S1 state produced by D.

Call this the *response-dependent account.* The idea is that pains and smell qualities are indeed response-independent physical properties of bodily regions or odor clouds. But the right account of their structural features is a response-dependent account, which brings in internal neural states. This basic idea has been pursued by Shoemaker, Lewis and McLaughlin (although they are certainly not tracking intentionalists!).[23]

But the problem with this reply on behalf of the tracking intentionalist is obvious: the response-dependent account is incompatible with tracking intentionalism (which goes with the fact that it is only non-tracking intentionalists who have endorsed it). As I said in the first premise of my arguments, on tracking intentionalism, phenomenal facts are facts about the properties represented by experience, *not* facts about the neural content carriers. To make the conflict vivid, just consider Harold and Henry from §4.2. The response-dependent account entails that Harold's first two smell experiences resemble more than his second two. But it entails that it is *not* the case that Henry's first two smell experiences resemble more than his second two, because of the difference in his responses to the molecules, that is, in his ensemble activations. This is incompatible with tracking intentionalism, because tracking intentionalism entails that Harold and Henry have exactly the same smell experiences, because they optimally track the very same chemical properties.[24]

---

[23] For a critical discussion, see Pautz 2006.

[24] For the same reason, an account of resemblances among the chemical properties with which the tracking intentionalist identifies smell qualities in terms of their effects on receptors (so that they're similar insofar as they produce similar receptoral activation patterns) is inconsistent with tracking intentionalism. (Just imagine a coincidental variation case in which the tracked chemical properties are the same but the receptoral patterns are different.) (There is also an empirical problem: resemblances among receptoral patterns doesn't need to match resemblances among smell experiences; it is only further downstream in the brain that we find such a match. Compare color vision.)

Jonathan Cohen made (in discussion) a point that I think is worth addressing: his worry was that the argument I give in the text for thinking that the response-dependent account is incompatible with tracking intentionalism is basically just the internal-dependence argument I previously gave

5.4 An Unsatisfactory Response-Independent Account

Now let me address another response the tracking intentionalist might offer (due to Michael Tye, in discussion). The response invokes what I will call the *unnatural relations* account of structure judgments. The previously considered response-dependent account of structure judgments was inconsistent with tracking intentionalism. This suggests that what the tracking intentionalist needs is a more response-independent account. The unnatural relations account is just such an account. The basic strategy is this: we explain structure judgments, not in terms of our neural responses to external properties, but in terms of the very unnatural response-independent correlates of those responses.

Let's first consider case where Max has the two pains. The unnatural relation response has two parts. First, if we conceive of relations abundantly, there is bound to be response-independent relation, R, among bodily types with the following property: if bodily disturbances D2 and D1 normally produce pain experiences E2 and E1, and subjects judge E2 to be roughly twice greater than E1, then D2 bears R to D1. Notice that we theorists just picked out the relation by mentioning our responses; but the idea is that the relation picked out is a wholly response-independent relation. This response-independent relation R is a very unnatural relation in David Lewis's sense. It is an uncodifiable disjunctive relation: roughly, *D2* and *D1* stand in *R* iff *D2* and *D1* are electric shocks applied to the skin with so-and-so spatial extents and *D2* is 25% greater than *D1*, *or D2* and *D1* are noxious temperatures with such-and-such spatial extents and *D2* is 33% greater than *D1*, *or* . . .But there is such a relation R. (Indeed: there are infinitely many of them!)

The second part of the unnatural relation reply is a semantic claim. It is that, when Max says that his right arm pain is twice greater than his left arm pain, the predicate 'x is twice greater than y' stands for the extremely unnatural relation U. This seems to require that this predicate is *ambiguous.* So on tracking intentionalism his report is true just in case D2 and D1 stand in the extremely unnatural relation U. Since they do, his report comes out true, as desired.

Now let us see how the unnatural relations account works in the case of Harold. Again, the account has two parts. First, there is an extremely unnatural four-place relation R such that A, Z, C and C* (in that order) stand in R. In general, this relation is the objective, physical correlate of judgments to the effect that two smells are more similar than two other smells. It is totally unclear what this relation would be like. Psychophysics certainly has no uncovered such a relation. But if we suppose that every miscellaneous class of ordered n-tuples determines a relation, then there certainly is such a relation. Second, in the context in which Harold makes his introspective report [S], his use of the predicate 'x and y resemble more than w and z' picks out R, and [S] is true iff A, Z, C and C* (in that order)

---

against tracking intentionalism; so that my structure argument requires the internal-dependence argument and consequently is not independent as I have said. This is not so. The conclusion of the argument I gave in the text is merely that the response-dependent account and tracking intentionalism are incompatible; the conclusion of the internal-dependence argument is that tracking intentionalism is false. So they are not the same argument. Moreover, the internal-dependence argument requires that in some coincidental variation cases the right verdict is actually Different Experiences. By contrast, my argument for thinking that the response-dependent account and tracking intentionalism are incompatible only requires that the response-dependent account entails Different Experiences in some such cases, as against tracking intentionalism.

stand in R. Since they do stand in R, his introspective report comes out true, as desired.

Notice that unnatural relations account, unlike the response-dependent account looked at above, is perfectly consistent with tracking intentionalism. The reason is that it is a response-independent account. Consider again Harold and Henry. In the counterfactual situation involving Henry, A, Z, C and C* (in that order) stand in R, just as they do in the actual world. So, on the unnatural relations reply, Henry's first two smell experiences resemble more than his second two, just as Harold's first two smell experiences resemble more than his second two. So this account is quite consistent with the verdict of tracking intentionalism that there are no phenomenal differences between Harold and Henry's smell experiences.

I think there are many problems with the unnatural relation reply. Here I will mention two.[25]

My first objection is this. The proponent of the unnatural relation reply says that, when Max makes his report [P], 'is twice greater than' in his mouth refers to the extremely unnatural relation U. He needs an account – a metasemantics – that explains how this is. I do not think that there is any credible metasemantic theory that might do that.

Of course, the general problem of metasemantics – how does language hook up to reality? – is everyone's problem. But I think that the proponent of the unnatural relation reply faces a special problem, indeed one that cannot be solved. For instance, we all have to say how 'is water' comes to pick out the property of being H20. But I can kind of see how that might be. For instance: the property of being H20 normally causes our use of 'is water' and our water-thoughts. By contrast, I cannot see how 'is twice greater than' in Max's mouth refers to the extremely unnatural relation U, as the proponent of the unnatural relations reply alleges. As we just saw, one kind of metasemantic theory appeals to causation. But the instantiation of the highly disjunctive relation U does not do a good job of causally explaining Max's use of 'twice greater than'. (The cause essentially involves internal neural factors, e. g. S1 firing rates.) Even if it does, then countless other such disjunctive relations to do – in which case we'd get the kind of indeterminacy worry discussed below. So the tracking intentionalist who favors the unnatural relations account cannot avail himself of a causal metasemantic theory in order to answer the metasemantic problem. There are other metasemantic theories that appeal to "naturalness" due to David Lewis and very popular among metaphysicians. The *very rough* idea is that a predicate refers to the most natural property or relation that fits our use of the predicate. But is the egregiously unnatural relation U the most natural relation that fits Max's use of the predicate 'is twice greater than'? Certainly not.

[25] There are other serious problems. One was suggested to me by Hans Kamp. When subjects in psychophysical experiments (e. g. Price, Myles) are asked to say when one pain is 'twice greater' than another in intensity, the understand the task immediately. This suggests that they are using their understanding of 'twice greater' in the usual way. This goes against the unnatural relation reply. On this reply, in connection with pain, 'twice greater than' takes on a special new semantic value (U) that it does not have in other contexts. Further: if 'twice greater than' really takes on a new semantic value in this context, what explains this? Not speakers intentions! (This adds to the metasemantic problem I am about to raise.) Also: there is no evidence at all that 'is twice greater than' is ambiguous in the required way.

Similar considerations apply to the unnatural relation account of Harold's report [S]. In fact, it is just absurd on the fact that his use of 'x and y resemble more than w and z' picks out any such four-place response-independent relation as R such that A, Z, C and C* (in that order) stand in that relation. What is the *response-independent* sense in which A and Z are more similar than C and C*? To appreciate the problem, look back at the picture in §5.1. Even putting this problem to the side, the view is totally unbelievable for metasemantic reasons. What could make it the case that Max's use of 'x and y resemble more than w and z' comes to pick out R in the relevant context? Sometimes context and speaker's intentions helps to determine what 'x and y resemble more than w and z' picks out in a context. But obviously speaker intentions will not help here. Nor will causation or naturalness.

The fact is, given "good internal correlation" and "bad external correlation", on any credible metasemantic story, if our structure judgments come out true, they concern internal factors, in a way that is incompatible with tracking intentionalism. For instance, one possible view is old-fashioned type-type identity theory. On this theory, Max's pains are not identical with the external bodily disturbances D2 and D1; they are identical with somatosensory neural states in the cortex. When Max makes report [P] that his right arm pain is twice greater than his left arm pain intensity, he's using 'twice greater than' to mean what it usually means, and he's correctly applying it to the neural firing rates that are involved in those neural states. This account of the truth-conditions of Max's report does not bring in any super unnatural relation such as U. So it beats by a long shot the account offered by the unnatural relation reply on the score of naturalness. Likewise, on the identity theory, smell experiences are ensemble activations; and Harold's report [S] is true just in case his first two ensemble activation states resemble more than his second two. Given the empirical evidence for "good internal correlation" (e. g. the Howard study), this is straightforwardly true. So on this account, by contrast to the unnatural relation account, there is no need to bring in any such unnatural relation as R. And this account jives with the fact it is resembles among ensemble activation states that causally explain Harold's report. (To explain our judgments about structural relations among odors in the world, the identity theorist might avail himself of the response-dependent account discussed in §5.3 above.) So, whereas I cannot see how the unnatural relation account might be true, I can see how *these* semantic accounts *might* be true: there are metasemantic theories (appealing to causation or naturalness) that might support them. But of course these semantic accounts require commitment to the identity theory (or some version of functionalism that similarly appeals to internal factors), and so are incompatible with tracking intentionalism.

So much for my first, metasemantic objection to the unnatural relation reply. My second objection is that, even if the metasemantic objection can be overcome, the unnatural relation reply entails that there is radical indeterminacy in certain of our structure judgments. To illustrate the objection, I will focus on the unnatural relation account of Harold's judgment. Of course, there is not just one response-independent relation R such that if chemical properties X, Y, W and Z stand in R then we judge the smell experiences produced by X and Y to be more similar to those produced by W and Z. There are infinitely many such relations: R1, R2, R3, ... They have the same extension when it comes to chemical properties of the kind to which our olfactory systems are causally responsive, but diverge

when it comes to other chemical properties. (Compare plus and quus.) Now what could make it the case that, when we discuss resemblances among smells, we are referring to one of these to the exclusion of others. Not *use:* their extensions differ when it comes to chemical properties that we don't encounter. Not causation: it is not as if one of these candidates to the exclusion of all others is uniquely the cause of our use of the resemblance predicate in connection with smells and smell experiences. Not naturalness: R1, R2, R3, . . . are, we may suppose, equally unnatural. (Contrast plus and quus.) So it looks as if the proponent of the unnatural relations reply would have to say that when we discuss resemblances among smells or smell experiences, it is radically indeterminate which of these relations we are referring to. But this is absurd. To see this, suppose we encounter a creature, Blurg, that has an olfactory system that is causally responsive to chemical properties that our olfactory system is not responsive to. Blurg maybe has alien olfactory experiences, just as a pigeon might have alien color experiences. Suppose Blurg has several smell experiences. And suppose we conjecture on the basis of his behavior that 'his first two smell experiences are more similar than his second two'. Surely this might be determinately true (or false), even if we cannot know this for sure. But on the indeterminacy response all such claims are indeterminate because whether such a claim is true hangs on whether we were referring to R1 or R2 or R3, and on the present view that is indeterminate! Notice that the identity theory semantics discussed above avoids this problem. For on this view Blurg's smell experiences are identical with his ensemble activation states; and there is a fact of the matter about whether Blurg's first two ensemble activations are overall more similar than his second two, even if we cannot access this fact on the basis of his behavior.

So, the basic point of the structure argument is that *there is no credible version of tracking intentionalism on which our structure judgments come out true.*

5.5 What's New in the Structure Argument?

There is a well-known argument against the view that colors are reflectance properties that concerns the phenomenal structure of the colors. Hardin said, in effect, that if this view is true, then our beliefs about the *unitary-binary structure* of colors are false. Byrne and Hilbert went beyond Hardin: they argue that, if this view is true, it *looks* as if our beliefs about *color resemblance* are also false.[26] Hardin endorses this sort of argument whereas Byrne and Hilbert reject it.

My argument is different from this argument and adds a number of fresh points.

(i) My argument is of course directed against tracking intentionalism about experience, not just a view about colors. It is not in the first instance about sensible properties like colors. It is in the first instance about experiences. It is meant to show that there is no credible version of tracking intentionalism on which our structure judgments about our *experiences* come out true.

(ii) One reply to the Hardin-Byrne-Hilbert argument about color structure appeals to the so-called "hue-magnitude" account. Byrne and Hilbert defend one version and Tye defends another. Now I myself think that this account is unsatisfactory even in the color case.[27] But it in any case simply does not carry over to the

[26] Hardin 1988; Byrne and Hilbert 1997 section 3.2 and foonote 32.

[27] Although in other respects Byrne and Hilbert are staunch supporters of a response-independent account of color, they somewhat surprisingly defend a response-dependent, *perceptual account* of

examples I have discussed involving smell experiences and pain experiences.[28] The tracking intentionalist would need some different account here. It is impossible to see what that account might be.

(iii) Another natural account of structure judgments about pain or smell brings in internal factors (the response-dependent account in §5.3). On this view, sensible properties are response-independent physical properties but claims about their structure are given a response-dependent analysis. This sort of account has been defended in the case of color by Shoemaker, Lewis and McLaughlin. I have shown that this account is not available to tracking intentionalists.

(iv) Finally, I have addressed an account not discussed in the literature (but mentioned by Tye in discussion), the unnatural relations account. This is the only account I can think of consistent with the tracking intentionalist's radically externalist account of phenomenal character. I have argued that this account fails on broadly metasemantic grounds.

(v) More generally, I have used "good internal correlation" and "bad external correlation", together with metasemantic considerations, to clinch the structure argument against tracking intentionalism: these considerations show that, at least if we assume some kind of reductive physicalism (no spooky, super-natural primitive sensible properties out there), the only credible accounts of our structure judgments on which they come out true are accounts that appeal to internal factors (e. g. neural states) in a way that is incompatible with tracking intentionalism.

---

resemblance structure in terms of "hue-magnitudes". (2003, 15). (As Byrne 2003 note 38 points out, Tye advocates an response-independent version of the hue-magnitude account – one which Byrne argue against, in favor of his response-dependent, perceptualist version.) For instance, consider the fact that purple is more similar to blue than to yellow. Byrne and Hilbert would say that this amounts to a *perceptual* fact, namely, the fact there is a hue-magnitude, namely *being bluish*, that all blue-appearing objects and purple-appearing objects, but not all yellow-appearing objects, (necessarily) appear to possess. Byrne (2003, 660 and note 38) would provide a similar perceptual account of the fact that purple$_{21}$ is a perceptual mixture of red and blue. (Byrne and Hilbert are clearly giving an account of what it is for colors to actually resemble and to have a unitary-binary character, not merely an account of what it is for them to appear to do so. This was confirmed by Byrne in correspondence but disputed by Hilbert in discussion; see also Byrne 2003, 660 and note 38.) As I said, I believe that the perceptual account is unsatisfactory. For instance, Byrne and Hilbert's perceptual account of comparative resemblance in terms of a difference in apparent hue-magnitude is clearly inapplicable when we are dealing with three *shades of blue* whose instances *all* appear to have the *same* hue-magnitude, *bluish*. So it does not provide a general account of the relation *color x resembles color y more than color z*. Further, the perceptual account is open to numerous counterexamples (Pautz 2003). (Byrne and Hilbert (2003, 55) reply to my counterexamples to their account of the unitary-binary distinction by claiming that they were after all not offering an account of ordinary claims of the form 'purple$_{21}$ is *reddish-bluish*', which they surprisingly regard as false because they insist it is only colored *objects* and never *colors* that are reddish or bluish. However, Byrne (2003, 660) *does* give a perceptual account of claims of a very similar form, such as 'purple$_{21}$ is *a perceptual mixture of red and blue*'. My counterexamples can be repeated *mutatis mutandis* against this account.)

In any case, the crucial point for us is this: even if the perceptualist-hue-magnitude account is workable when it comes to claims about colors, nothing like it apply to claims about the phenomenal structure of pain experiences or smell experiences, such as Max's claim [P] and Harold's claim [S].

[28] Byrne and Hilbert (in 'Sensible Qualities and The Structure of Appearance') do have an interesting discussion of sensible properties other than colors but they do not say how the intentionalist who accepts a response-independent, physicalist account of smell qualities (parallel to their account of color) might accommodate the likes of Harold's report [S].

6 Third Argument: The Percipi Argument

My previous two arguments against tracking intentionalism were empirical. My third and final argument is that at least in the case of pain it looks like tracking intentionalism can also be ruled out on apriori grounds.

Let's begin with color, where I think that the apriori argument I am about to give does *not* work. Suppose you focus on a color. On tracking intentionalism, that is a mind-independent reflectance property. It could have occurred – it could have been instantiated – even if no one were around who is the subject of a color experience. This jives with our intuitions on the matter.

Now on tracking intentionalism, the metaphysics of pain experience is basically just like the metaphysics of color experience. If you have a pain in your leg, and focus on the horrible quality ostensible present there, you are in fact focusing on a mind-independent type of bodily disturbance. The trouble is that in this case, by contrast to the color case, we have a kind of percipi intuition which seems to show that tracking intentionalism just does not make sense in the case to pain.

To illustrate, suppose we're discussing a severe pain felt in your arm: I say:

Percipi Intuition: *Necessarily, if it occurs, then some one has a pain-experience*

Intuitively, my statement is true. But, on tracking intentionalism, my statement is false, since 'it' here refers to a type of bodily disturbance D could occur without a pain experience, for instance in an insentient cadaver. But my statement is true, so tracking intentionalism is false.[29]

A clarification: My claim – my percipi intuition - is that "it" cannot occur without an experiencer having a pain. A counterexample would be a case of the pain-quality without anyone experiencer around having a pain. So he case of non-focal pain is *not* a counterexample. That is a case of an experiencer having a pain who is not attending to his pain.

Let me address some responses.

(i) The first response is due to Michael Tye. Michael Tye appears to anticipate the percipi argument. (A similar reply might be suggested by Dretskes' 2003 discussion of crocks and protopain.) He writes:

> we apply the term 'pain' to tissue damage only in a certain context -- the context provided by tissue damage being represented by a [certain pain experience]" [this is why] "in a world without experiencers, there can be no pain but there certainly can be tissue damage" (2006)

On one interpretation, what Tye is saying here is that 'pain', when used to describe the property represented, is short for a description, 'the property represented by so and so pain experience'.

[29] Stysma (consciousness online) has a paper in which he presents "experimental" evidence that might seem to suggest that the folk don't have the percipi intuition, but as Chalmers points out (in the comments) the experimental design seems to be bad and the vignette clearly stacks the cards.

This descriptivist maneuver may validate some of our intuitions. Suppose I say:

> [1] The pain in my foot couldn't occur without an experience.

On the relevant descriptivist account, this becomes:

> [2] The pain represented by my pain experience as in my foot couldn't occur without a pain experience.

And [2] is true relative to a "de dicto" reading that gives the description narrow scope. The percipi intuition, so the idea goes, expresses a de dicto necessity, not a de re one.

But, by design, the descriptivist maneuver does not validate the initial intuition. Suppose again that we are discuss a sever pain in your arm. I say: Necessarily, if *it* occurs, then some one has a pain-experience. This sentence does not contain 'pain'. Even if some occurrences of 'pain' are abbreviated descriptions, the descriptivist, de dicto maneuver does not apply here. This sentence is unamiguiously *de re*, and if tracking intentionalism is true, then it is false, since here 'it' refers to a type of bodily disturbance which might occur without anyone having a pain experience. But it is true. So tracking intentionalism is false. So Tye's response to the percipi argument seems to fail.

(ii) It might be said (and was said by William Lycan in discussion) that the tracking intentionalist can accept the percipi intuition after all. The idea is that our talk of pain is ambiguous: sometimes it refers to bodily disturbances, sometimes it refers to neural events in the brain, token-experiences. Further, the idea - which is different from Tye's idea above – is that when I said "Necessarily, if it occurs, then some one has a pain-experience", my use of 'it' didn't refer to D, rather it referred to a neural event in the brain, an experience-token. Further, the response continues, this neural event is *essentially* an experience. (Notice that on tracking intentionalism, which is radically externalist, this is not obvious; but I will let this slide.) In this way, my modal claim can come out true, even on tracking intentionalism. (Indeed, the tracking intentionalist might even say that our talk of pain is three ways ambiguous: it might refer to (i) the properties represented by our pain experiences, (ii) the pain experiences in the brain, or (iii) the pain experiences plus a higher-order thought to the effect that one has the pain-experiences.)

The trouble with this reply is as follows. To begin with, I find this ambiguity view implausible; it is not motivated by any of the usual semantic tests. But let's just grant it for the sake of argument. Surely the tracking intentionalist must allow that some contexts when we are discussing pains we are discussing the qualities represented by our experiences, not merely the events in the brain that represent those qualities. Consider again that his account of pain experience is perfectly parallel to his account of color experience. When it comes to color experience, the tracking intentionalist surely allows that we can focus on and discuss the colors represented by our color experiences, not just the color experiences. And when we do we focus on qualities that occur without experiencers around. So, by parity, he must allow that when it comes to pain experience, when can focus on and discuss the qualities represented by our experiences, not just the experiences

themselves. Indeed, since he advocates "transparency", he might say that this is the usual case. The present response allows this by saying that our talk of pain is ambiguous. Now, in the context in which I said, 'Necessarily, if it occurs, then some one has a pain-experience', the contextual circumstances were meant to be such that on tracking intentionalism I could not be referring to anything but the quality represented by my experience, just as in the color case: we were talking about a pain in your foot, your attention was focused out, not in, and so on. So on tracking intentionalism I was referring to D, not some internal neural event, which could occur without any experiencer around. So tracking intentionalism violates the percipi intuition.

(iii) Finally, I think Chris Hill would just reject the percipi intuition that I have in mind. The blow would be softened if he could explain why we have this intuition. Hill (2006) *appears* to provide the following explanation:

1 Our concept of pain tracks neural-representations (pain experiences in the head), which evidently entail pain-experiences.
2 So it becomes part of our concept that pains entail pain experiences
3 Yet our concept of pain in fact refers to represented pain qualities, which are bodily disturbances; so in fact pain doesn't entail pain experiences.

But this explanation is psychologically implausible. The idea is that when I have a pain in my foot and have the percipi intuition, I am thinking about the something that's not in my head: the quality represented. This thing, according to Hill, can occur even if no one is having a pain experience. But because I'm often interested in some other thing that is in my head (neural-representation) and that does entail pain-experiences, I think this thing outside my head entails pain-experiences. Am I really that confused?[30]

7. Objections and Replies

I have laid out my three independent arguments against tracking intentionalism: the internal-dependence argument, the structure argument and the percipi argument. I have already considered some objections along the way. But there are yet other objections that need to be addressed.

First objection: impossible (Byrne and Tye 2006). Here is an objection directed only at the internal-dependence argument. The first premise was that on tracking intentionalism the correct verdict in all coincidental variation cases is Same Experiences. The second was that in some coincidental variation cases the right verdict is instead Different Experiences. But the second premise cannot be right. It is part of the stipulated conditions of coincidental variation cases that the individuals involved track the same properties under optimal conditions, in the sense of optimal conditions invoked in Tye's psychosemantics. It is incompatible with this

[30] Hill's present view (in his book *Consciousness*) appears to be a no-reference view according to which 'pain' just doesn't refer. On this view, I suppose all pain discourse is literally false. But even the proponent of the no reference view must explain away the intuition in favor of the percipi intuition (not to mention explaining away our error in saying things like 'I have a pain in my foot').

that they should have different experiences. If they have different experiences, then they represent the same items as having different smell qualities, taste qualities or pain qualities. Those qualities exclude each other. So if they have different experiences, then at least one must be constantly misrepresenting the smells, tastes or pain-qualities of external items. This is inconsistent with the optimality assumption, for part of optimality is *getting it right.*

Reply: This objection fails to answer my arguments for three reasons. First, the error assumption here is implausible: it is implausible to think that, if the individuals involved have different smell experiences, taste experiences or pain experiences, then one must be systematically misrepresenting the world. Second, even if we grant this, it is not incompatible with the optimality assumption, for optimality cannot be defined in terms of getting it right. Since it *figures in what it is meant to be a reductive theory of experience and its intentionality, it must be defined in non-intentional terms*.[31]So error is quite compatible with optimality in the relevant sense. So even if we accept the error assumption, we should accept Different Experiences, for all the reasons I gave earlier – in which case tracking intentionalism is false. Third, this objection does nothing to answer the structure argument or the percipi argument.

Second objection: seems like this (Dretske in discussion). If Yuck and Yum, Mild and Severe, or Harold and Henry are asked with it is like for them, they will say it smells, tastes, feels like this. Here 'this' refers to the response-independent, physical property they both track. So they will say that their experiences are phenomenally identical. So we have an argument against Different Experiences; and in favor of Same Experiences, which is in line with tracking intentionalism.

Reply. First of all, Yuck and Yum, Mild and Severe, and Harold and Henry may not speak a language like English. Even if they do it is up for grabs what their utterances mean, just as it is up for grabs what their experiences are like. It is like a case of radical interpretation. So we cannot make assumptions about what their words mean and use them to derive verdicts about what their experiences are like, in the manner proposed by the objector. The issue of what they mean when they describe their experiences, and what their experiences are like, should be decided together. I have argued that the most reasonable view is that they have different experiences, so that even if both say 'it seems like this', by 'this' they are referring to different things. The objection does not answer those arguments.

In addition, the objection does not answer the structure argument or the percipi argument.

Third objection: just a difference in affect/desire/valence (Dretske, Lycan, in discussion). Granted tracking intentionalism entails that Yuck and Yum and Mild and Severe have exactly the same sensory experiences of the berries and the armlesion (respectively), because they optimally track the very same external properties. But, at least when it comes to these cases, that is acceptable: the neural and behavioral arguments for thinking that they have different sensory experiences fail. True, they differ markedly in their neural processing, but this does not mean

[31] For this point, see sect. 5 of https://webspace.utexas.edu/arp424/www/simple2.pdf, Chalmers https://webspace.utexas.edu/arp424/www/simple2.pdf, Cohen The Red and the Real, p. 83 note 46.

that their experiences differ in sensory phenomenology. I just choose reject the empirically-based arguments offered for saying that their experiences differ in sensory phenomenology. True, they differ in behavior. But this doesn't show that they have experiences that differ in sensory phenomenology. Maybe Yuck and Yum have exactly the very same sweet taste experience (for example); but, whereas Yum and everything else in his species likes this sweet taste experience, Yuck and everyone else in his species really hates it, which explains why they withdraw and so on. Likewise, maybe Mild and Severe have exactly the same pain experience, but it just bothers Severe more, which is why he responds with greater urgency than Mild. And maybe these differences in desire or affect are grounded in the functional differences between them.

Reply. I have three points in reply.

First, my main reply is that, even if this objection were successful in the case of Yuck and Yum and Mild and Severe (which I deny for reasons I will explain presently), it does not carry over to other coincidental variation cases. Consider, for instance, the case of Harold and Henry. We can suppose that both Harold and Henry behave as if they like the four odors they smell equally well. So their experiences have the same valence ("good"): there are not differences in valance. A fortiori, we cannot explain the fine-grained behavioral differences between them (in categorization and discrimination) just by positing differences in valence. Given the vast neural and behavioral differences between them, and given "good internal correlation" and "bad external correlation" in the domain of smell, the only reasonable view is that their experiences differ at sensory level (viz. they differ in their resemblance structure). The "it's just an affective difference" objection also does not carry over to other coincidental variation cases I have described elsewhere concerning color vision and sound.[32] These kinds of experiences do not even have a significant affective component. So in these cases the idea that the behavioral (ordering, sorting) differences between the individuals involved are explained by the assumption they have the same color and sound experiences but differ in their affective responses to them is even more of a non-starter. Now, as I stressed earlier, to refute tracking intentionalism, there only needs to be one possible coincidental variation cases in which the right verdict is Different Experiences but tracking intentionalism delivers the verdict of Same Experiences. So, even if the 'it's only an affective difference' objection were successful in the cases of Yuck and Yum and Mild and Severe, the internal-dependence argument against tracking intentionalism would be untouched.

Second, the 'it's only an affective difference' objection is in any case unsuccessful even in the cases of Yuck and Yum ad Mild and Severe. The neural and behavioral differences between Yuck and Yum, and Mild and Severe, make it reasonable to suppose that they have experiences that differ in *sensory* phenomenology. The sorts of neural differences that exist between Mild and Severe, and between Yuck and Yum, are ones that we know to correlate well with differences in *sensory* phenomenology. In particular, Mild and Severe differ in S1 firing rates. And S1 firing rates are known to be well correlated pain intensity at the *sensory level*. Lesions within S1 produce deficits in ability to detect sensory features of pain in both monkeys and humans. S1 activity is finely graded in response to nox-

[32] Pautz 'Sensory Awareness is not a Wide Physical Relation' and 'Do Theories of Consciousness Rest on a Mistake?'.

ious stimuli, and resembles the response characteristics of S1 neurons parallel psychophysical characteristics of pain in human and non-human primates. In addition, Hofbauer (2001) and Rainville associates have shown that if hypnosis is used to reduce the unpleasantness of pain while keeping the sensory intensity of pain the same, S1 activity also remains the same. What changes is activity in the anterior cingulate cortex (ACC), showing that S1 activity is responsible for the sensory aspects of pain and ACC activity is responsible for the affective aspects of pain. So, for all reasons, the fact that Mild and Severe differ in S1 activity is very good evidence that they differ in pain at the sensory level, not just in the affective level – not just at the level of desire or affect. Similarly, the sorts of differences that exist between Yuck and Yum are ones that we know ground differences in sensory phenomenology. Salty food and sweet foods taste good – experiences of them are alike at the affective level - but they are qualitatively different. This qualitative difference is matched by a difference in their ensemble activation patterns in the brain. So it really looks as if across fibber patterns are responsible for the sensory dimension of taste. Since Yuck and Yum differ in their across-fiber pattern states, it is very reasonable to suppose that they taste experiences that differ systematically at the sensory level, not just at the affective level. It is not the case that both have a sweet taste. Rather, maybe Yum has a sweet taste and Yuck has a bitter one. In addition, I have supposed that Yuck and Yum exhibit *innate* and very *fine-grained* differences in behavior in sorting behavior, differences which are widespread in their species: ordering behavior, and so on. These sort of fine-grained differences are not well-explained by saying that they just differ at the affective level; they are only well-explained by saying that their differ at the level of sensory phenomenology. (Here I'm indebted to Fred Dretske.)

Third, the 'it's only an affective difference' objection to my internal-dependence argument does nothing to answer the structure argument and the percipi argument against tracking intentionalism, which are entirely independent of the internal-dependence argument.

Fourth objection: peripheral neural activity (Byrne, Hill, in discussion; also Hill forthcoming) The first premise of the internal-dependence argument is that on tracking intentionalism the correct verdict in every coincidental variation case is Same Experiences. The second premise is that in some such cases the correct verdict is rather Different Experiences. The previous three objections questioned the second premise. But we can also question the first.

Granted, maybe the version of tracking intentionalism considered delivers the mistaken verdict of Same Experiences. But consider a view in the vicinity of tracking intentionalism: we might call it *peripheral tracking intentionalism*. The idea here is that represented (perceived) smell qualities and taste qualities are not chemical properties. Rather, they are properties involving the *peripheral neural activity* that those chemical properties cause. (Or maybe properties involving activity further downstream, even in chorda tympani nerve.) Likewise, represented pain qualities are not types of bodily disturbance; rather they are properties involving peripheral neural activity (C-fiber, A-fiber) caused by bodily disturbance. (Or maybe they properties involving activity further downstream, e. g. in the spinal cord - a suggestion of Byrne and Hill.) Now, suppose that Mild and Severe, Yuck and Yum and Harold and Henry differ in their peripheral neural activity. Then their experiences might track and hence represent different properties

involving that activity. So peripheral tracking intentionalism delivers the verdict of Different Experiences in coincidental variation cases, and the internal-dependence argument collapses.

Indeed, peripheral tracking intentionalism might also be less vulnerable to the structure argument than the kind of tracking intentionalism so far considered. For maybe the resemblance structure of our taste and smell experiences is matched by the resemblance structure of the peripheral activity. And maybe, when Max's second pain is twice greater than his first, the level of represented peripheral activity is roughly twice greater. The peripheral view is also natural in the case of thermoperception. For one thing, the difference between hot and cold qualities seems to be more than a difference in degree: there's a differences in kind. This sits ill with the view that these qualities reduce to high and low temperatures. On peripheralism, they would be identified with properties involving the activity of quite different receptors, the warm and cold receptors, in the skin.

Reply. Peripheral tracking intentionalism is pretty hopeless, I think. Contrary to the objection, it is just as open to the internal-dependence argument and the structure argument as the kind of tracking intentionalism I have discussed. It is also open to the percipi argument. And it is open to two additional, decisive objections. In all, it faces four objections.

First, peripheralism *is* open to the internal-dependence argument. For just suppose that Mild and Severe not only track the same bodily disturbance; that bodily disturbance also sets up the same peripheral neural activity. The differences between them in S1 firing rates are due solely to *cortical* differences. Then peripheralism entails that they represent the same peripheral neural properties and hence have the same experiences. But this is not at all plausible, given the downstream neural and behavioral differences between them. This case is not far fetched. There is reason to believe that there are cases like this: John and Jane response to the same bodily disturbance, but John has a much more severe pain than Jane, even though conditions are totally normal (no drugs, top down influence, etc.). Coghill and coworkers (2003) found that in such cases there are also differences between John and Jane in their levels of cortical neural activity. But interestingly they only found cortical differences; it appeared that peripheral, spinal, and thalamic activity was the same. As they put it, "generally similar afferent input was conveyed to thalamic levels in both high- and low-sensitivity individuals" (8542). The case of Mild and Severe, then, is similar to actual cases: the only differences are (i) I stipulate that they belong to different species and that the neural and behavioral differences are greater and (ii) I stipulate that under optimal conditions their brain states track *exactly the same* range of external bodily properties (which might be dispute in actual cases like John and Jane). Similarly, we can suppose that Yuck and Yum and Harold and Henry undergo the same peripheral activity: the neural differences are entirely postreceptoral. I have described a case like that involving color vision elsewhere (Max and Twin Max), in which the receptor activity is the same and the neural differences are further downstream. In such all such cases, even peripheral tracking intentionalism delivers the mistaken verdict of Same Experiences.

Second, peripheral tracking intentionalism is open to the structure argument as illustrated by the cases of Harold and Max and other such cases. There is no reason to think that the resemblance structure of our taste and smell experiences

is matched by the resemblance structure of the peripheral activity. Compare color vision: the unitary-binary structure and resemblance structure of color experience has no echo at the level of the receptors. In the case of taste and smell, as in the case of color vision, it is only further downstream in the brain that we find a "match". In the case of pain, peripheral neural activity poorly reflects pain intensity, because of downstream temporal and spatial summation. So peripheral tracking intentionalism, as well as the kind of ordinary tracking intentionalism I have mostly focused on, seems unable to accommodate the truth of our introspective structure judgments.

Third, of course, peripheral tracking intentionalism about pain is just as open to the percipi argument as the more usual form of tracking intentionalism on which represented pains are types of bodily disturbance.

Fourth, peripheral tracking intentionalism is open to objections of a more technical nature. So it is actually much *worse* than the more ordinary kind of tracking intentionalism I have considered. (i) A well-known problem in psychosemantics is the "distance" problem. There are many properties involved in causing a brain state: properties out in the world, properties involving peripheral neural activity in the receptors, properties of the nerves going from the receptors to the brain, and so on. What makes it the case that an internal state represents one of these properties to the exclusion of the others? (We surely do not represent all of them. Or at least we do not sensorily or phenomenally represent all of them.) The peripheralist must solve the distance problem in such a way as to uniquely favor the properties involving the receptors. This seems impossible. For instance, to solve the distance problem, one might invoke the notion of function. But surely, if anything, our brain states have the function of indicating biological important properties like level of sugar and bodily damage, not properties involving peripheral neural activity. So the teleological solution answers the distance problem in way inconsistent with peripheralism. (ii) Nevertheless, suppose for the sake of argument that the right psychosemantic theory solves the distance problem in a way that entails that our smell experiences, taste experiences and pain experiences represent nothing but properties involving peripheral neural activity. Then, by parity, it should entail that our visual experiences represent nothing but properties involving the activity of the receptors on the eye. Represented colors are not reflectances or the like but properties instantiated in the retina! This is absurd for many reasons. (iii) Versions of peripheralism on which pain experience represent properties involving spinal cord activity must explain how our experiences manage to mistakenly represent these properties as instantiated in various bodily reasons. This seems impossible.

Fifth objection: similarity (Byrne, in discussion). How can you be so sure that tracking intentionalism delivers the objectionable verdict that Harold and Henry have the same smell experience, that Mild and Severe have the same pain experiences, and so on? Don't you have to know first how the tracking intentionalist accounts for phenomenal structure in the domains of smell and pain and so on?

Reply. First, I have no idea how the tracking intentionalist might account for phenomenal structure in the domains of smell and pain and so on. (Nothing like Byrne and Hilbert's hue-magnitude account seems possible in these cases.) This is itself a problem, which the objection does not answer. That was the structure

argument. Second, suppose he develops some view on this issue: he has some account of what resemblances among smells come to, for instance. As Byrne argues (Color and Similarity, sect. 6), it is implausible that our experiences represent resemblance relations among sensible properties. (It is especially implausible when subjects are only presented with one sensible property, as we may suppose is the case with the individuals in coincidental variation cases!) So, the tracking intentionalist cannot say that the individuals in coincidental variation cases have different experiences because they represent different resemblance relations among sensible properties, for the simple reason that our experiences (as opposed to our beliefs) arguably do not represent such relations at all. Third, even if they do, the tracking intentionalist cannot say that the individuals in coincidental variation cases represent different resemblance relations, for by stipulation they are exactly alike in what they track, and on tracking intentionalism all aspects of representational content are determined by tracking. Tracking intentionalism is stuck with the verdict that their experiences have exactly the same representational contents, and hence exactly the same phenomenal characters. But, to repeat, this verdict is incredible.

Fifth objection: evaluative intentionalism (Hill in discussion and forthcoming). The reductive intentionalist might say that Yuck's experience represents the berries as bad and Yum's represents them as good, and that this constitutes the sensory-phenomenal difference in their taste experiences. For instance, this is why Yuck has a bitter experience and Yum has a sweet one.

Likewise, he might say that Mild's experience Severe's pain is experienced as more intense because the threat that it represents is more severe, despite being associated with the same lesion, given the set-up of the case. In short, Mild's experience represents it as being threatening to degree D1 and Severe's represents the same lesion as being threatening to degree D2.

Perhaps this threat-level idea could also help with the example of Max used in the structure argument. If Max is stimulated by two electric shocks, the second of which is greater but less than twice greater than the first, it may be that the lager shock poses a threat to the integrity of the organism that is roughly twice the threat that is posed by the first shock (Hill forthcoming). If Max's pain experiences represent these treat levels, the reductive intentionalist can accommodate his judgment that his second pain is roughly twice greater in intensity than his first.

So the intentionalist might accept evaluative intentionalism about some aspects of some experiences.

Reply. This objection is not really a defense of tracking intentionalism. Accepting evaluative intentionalism would require rejecting tracking intentionalism. On tracking intentionalism, sensory phenomenology (sweet vs. bitter, etc) is determined by the representation of response-independent properties, for instance chemical properties and bodily damage understood in response-independent terms. On evaluative intentionalism, by contrast, it is determined at least in part by the representation of properties like badness and harm-level. (Tye holds that the representation of badness can contribute to affective phenomenology but denies that it contributes to sensory phenomenology.) Also, given how I set up the case of Yuck and Yum, the tracking intentionalist cannot say that Yuck's experience represents the berries as being bad (or poisonous) at all. For the poisonous-

ness of the berries is grounded in chemical property, T, different from C, which is undetectable by the taste system and which only three-fourths of the berries possess. So on Tye's causal psychosemantics and on Dretske's indicator psychosemantics (which requires probability 1), Yuck's experience simply cannot represent this property.

So accepting evaluative intentionalism would involve rejecting tracking intentionalism. But let me consider it anyway. I think it is not promising for three reasons. (i) Contrary to evaluative intentionalism, fine-grained sensory phenomenology certainly cannot be grounded in the representation of coarse-grained properties like being good or being bad. For instance, if Yuck has two bites of the berries, and gets slightly different sweet taste experiences, both experiences would represent the berries as good. (ii) Evaluative intentionalism in any case cannot be enlisted to answer the internal-dependence argument in other cases. Elsewhere I have described coincidental variation cases involving color experience and sound experience. Here the phenomenal differences between the individuals involved certainly cannot be explained by saying that their color experiences or sound experiences differ in whether they represent things as good or bad: such experiences simply do not represent things as good or bad. So to handle the cases the reductive intentionalist would have to come up with some other theory (like the Kriegel-Prinz theory to be discussed in the next section). (iii) Evaluative intentionalism does nothing to answer the structure argument. On this view, what truth does Harold express, when he says that his first two smell experiences resemble each other more closely than his second two? Or again, on evaluative intentionalism, Max's report that his second pain is twice greater in intensity than his first pain is true because his second pain represents a threat level T2 and the second represents a threat level T1, and T2 is twice greater than T1. But this is not much progress until we are told what threat levels are, and just when one is "twice greater" than another. (iv) Whatever threat levels are, they are going to be highly relational properties which involve the overall condition of an organism. It is hard to see what psychosemantic theory of sensory representation might explain how our experiences represent such properties. Likewise for properties like *being bad*.

Sixth objection: (Hill in discussion) Maybe we should appeal to relations between inner neural states (e. g. ensemble activations) in explaining certain relationships of apparent similarity and difference between sensible properties. At this point, it isn't clear to how much of phenomenal character can be explained in terms of representational content, and how much is to be attributed to non-semantic properties of inner neural states.

Reply: This objector in effect suggests rejecting tracking intentionalism. So what is said here is not really a defense of tracking intentionalism against my arguments. It concedes that those arguments are successful against tracking intentionalism. Indeed, if we understand intentionalism in a strong way, so that it means 'the phenomenal character of an experience is the set of qualia that the experience *represents as instantiated by objects of awareness*' (Hill *Consciousness*, 148), what the objector suggests involves rejecting *intentionalism*. The idea seems to be that the Yuck and Yum, Harold and Henry, and Mild and Severe have experiences with exactly the same representational contents, but those experiences differ in

phenomenal character due to the neural (and perhaps functional) differences between them.

Incidentally, although I reject tracking intentionalism, I do not think that the view suggested in the objection is promising. Imagine having Yuck's taste experience and then Yum's. You'd would experience a numerically different taste quality, apparently in your tongue. Likewise, if you were to have Mild's pain experience and then Severe's, you would experience a numerically different pain quality, ostensibly located in your arm. On the view suggested by the objector, what physical properties are these qualities identified with? Neural properties in the brain? That would help to explain why Yuck and Yum, and Mild and Severe, perceive different such qualities – for different neural properties are instantiated in their brains. But of course an obvious problem with this suggestion is that it is hard to see what theory of intentionality could explain why we Yuck and Yum and Mild and Severe experience these qualities as located in or around their tongues and in their arms. Of course, in the domain of color vision, the idea seems even more problematic: the idea would be that sensible colors are neural properties, which we somehow experience as pasted on external objects. I also worry that the view suggested in the objection would just collapse into old-fashioned central-state materialism, with all of its problems.

8. The Kriegel-Prinz Theory

Now I have completed the main task of my paper. I have shown that tracking intentionalism does not work. Now I'd like to look at an alternative theory due to Uriah Kriegel (2009) and Jesse Prinz (2006). I'd like to look at it because it might have naturally occurred to the reader: one might think that is a natural view to take in view of my arguments against tracking intentionalism.

Like tracking intentionalism, the Kriegel-Prinz theory has three parts.

(i) The first part is intentionalism.

(ii) The second part is a response-*dependent* theory of the sensible properties represented by our experiences. It contrasts with the response-*independent* theory of sensible properties that is part of tracking intentionalism. On Kriegel's version, sensible properties are *infinite conjunctions!* They have the form: *normally causing brain state B individual X under conditions Y & normally causing B* in X* under Y* & . . .* Prinz says Prinz says that "sensations represent appearances" (2006, 451) and appearances "are powers that external things have to cause representations in our dedicated input systems" (438). But he does not say why he accepts response-dependent reductionism over response-independent reductionism, or select between the many versions of it. On a very simple version of this view, which I will be working with for convenience, tastes, pains, and smells are properties of the form: causing, or normally causing, brain state B, where B is a brain state that realizes a particular experience.

(iii) The third and final part of the Kriegel-Prinz theory is a theory of sensory representation along the lines suggested by Tye and Dretske in terms of causal covariation or indicator function. As it happens, both favor Dretske's psychosemantics. But, for completeness, I will also consider how their view meshes with Tye's optimal tracking psychosemantics.[33]

[33] I have attributed intentionalism to both Kriegel and Prinz. Kriegel is clearly some kind of inten-

In short, on the Kriegel-Prinz theory, intentionalism is true: sensory consciousness represents bodily regions, food-stuffs, and so on, as having various properties. But the sensory systems are self-centered: they represent extracranial items as having properties that involve how they affect the central nervous system. William Lycan briefly suggested a similar reply to Levine's fish case (although if what I said previously about this case is right the reductive intentionalist can give a much simpler reply, namely that the case is not possible). This view is already defended by Sydney Shoemaker. It is also defended, albeit more tentatively, by Gilbert Harman.[34] But there is a huge lacuna in Shoemaker's (and Harman's) theory. He has no theory of the "sensory representation relation" that we allegedly bear to the relevant response-independent properties. Kriegel and Prinz are basically taking Shoemaker's view but filling the lacuna by providing a broadly Dretskean theory of this relation on which this relation reduces to what I previously (in §2) called "the indication relation". This is what makes the view distinctive (and also what makes it unworkable, as we shall see).

Now one might think that the Kriegel-Prinz theory avoids my three arguments against tracking intentionalism. I think this is the only possible argument for the theory, although it is not one that Kriegel and Prinz themselves make.

For instance, one might think it avoids the argument about internal-dependence. Suppose that on tasting the berries Yuck undergoes ensemble activation A1 and Yum undergoes ensemble activation A2. Then maybe on the Kriegel-Prinz theory, Yuck and Yum don't optimally track, and thereby represent, just the response-independent chemical property C, possessed by the berries. Rather, Yuck optimally tracks, and thereby represents, the response-dependent properties of the berries of normally causing across-fiber pattern A1. And Yum. It is this representational difference that grounds the phenomenal difference among their taste experiences, one might think. In fact, if the intentionalist wants to accommodate internal-dependence, it almost seems inevitable.

But the Kriegel-Prinz theory faces many problems. Here I will focus what I take to be the main problems.

The first problem is the simplest and most obvious: I call it *the incompatibility problem*. Hill (2006) and Levine (2010 ND review of Kriegel) have also raised this problem; but I will add to what they say.

---

tionalist in his 2009 (Subjective Consciousness). But, despite having written a great deal about consciousness, Prinz has not as far as I know explicitly taken a stand on the intentionalism/anti-intentionalism debate. He has also focused on vision, mostly ignoring the other modalities. Still I tentatively attribute to him some kind of intentionalism. For instance, he writes "Sensations do represent those things that make a phenomenal difference, and those things are appearances" (2006, 458), which at least commits him to a kind of partial intentionalism. He also (2010, 313) accounts for qualitative character in the domain of vision in terms of "intermediate level visual representations". This suggests some kind of intentionalism about visual experience. (By the way, I think this account is problematic. "Intermediate level" is vague. And: does this notion have application to all possible subjects of visual experience? Arguably not. So the intermediate level hypothesis, if true, is a contingent truth about some creatures; it does not tell us *what it is to have a visual experience with a certain character*, which is usually what philosophers want from a physicalist or other theory of visual experience. Finally, of course it is hard to see how this intermediate level theory might extend to non-visual experiences.

But even if Prinz does not accept intentionalism (i), he does accept (ii) and (iii). This makes him vulnerable to the problem I will raise, for that problem is that (ii) and (iii) simply do not go together.

[34] Harman 1996.

The Kriegel-Prinz theory attempts to combine Tye or Dretske's view of sensory representation with the view that the represented properties are response-dependent properties of the form normally causing brain state B. The theory fails for the simple reason that these two views are mutually inconsistent. In short, it is the conjunction of (i)-(iii) but (ii) and (iii) are incompatible.

On Tye's view, a brain state B represents "the" external property that's "causally efficacious" in the production of B under optimal conditions, where "causal efficacy" is an alleged objective, context-independent notion. But the dispositional, response-dependent property *normally causing brain state B* isn't causally efficacious in the production of brain state B. The properties that are, are response-independent properties: chemical properties, bodily disturbances, etc. So on Tye's theory of representation, we sensorily represent the response-independent properties, not Kriegel & Prinz's response-dependent properties. That's of course Tye's view. Consider an analogy. A state S of a thermometer causally-covaries with, and thereby represents, a response-independent temp, not the response-dependent causing that very state, S.

Now let's turn to Dretske's view. On Dretske view, a brain state B belonging to a sensory system represents "the" external property that the brain state has the "function of indicating". *Presumably,* a brain state B does not have the "function of indicating" the biologically unimportant response-dependent property normally causing brain state B. A brain state has the function of indicating a biological important property: chemical property, bodily disturbance, etc. So on Dretske's theory of representation, we sensorily represent such response-independent properties, not Kriegel-Prinz's response-dependent properties. That is of course Dretske's view.[35]

Another point is worth making. In response to some of Stich's persuasive objections, Dretske writes: "when an [internal state], C, indicates both F and G, and its indication of G is via its indication of F, then C acquires the function of indicat-

[35] It might be worth mentioning that Dretske's theory of sensory representation is actually quite undeveloped. It differs from his account of how thought represents. In the case of thought-representation, he provides a partial, rough gloss on 'function to indicate': roughly, a state S has the function to indicate F only if S was recruited to lead to some action A because S indicates F. As Stich argues (1990), even with this gloss, the notion is quite vague and leads to indeterminacy, as Dretske admits in his response to Stich. Dretske's theory of *sensory* representation (which as far as I know he only develops at the start of *Naturalizing the Mind*) is more under-developed, in my opinion. In this case, he says, it is whole systems that have an indicator function – this is primary. (He says sensory representation is *systemic*, not acquired as Dretske thinks thought-representation is.) The indicator function of particular states (which I guess is generally to indicate determinate properties) is somehow derived from the function of the whole system (which I guess is to indicate a determinable like bodily disturbance). He also says – and this seems crucial to his account – that our brain states in the various sensory modalities have some kind of structure that matches the structure of properties in the world; and that this means our brain states can represent certain determinate properties even if there is no distinguishable evolutionary process whereby they explicitly acquired an indicator function. But this is very vague. What does he mean by the structure of our brain states? Does he really think that the structure of our brain states somehow matches the structure of properties in the world? Isn't this undermined by "bad external correlation"?

ing that F" (1990, 826). (Dretske then notes that this is very vague, so his theory is as yet incomplete.) Now our internal states indicate dispositional, response-dependent properties like being disposed to produce brain state B via indicating the chemical properties, bodily disturbances, and so on, which ground those dispositions. So, again, we see that Dretske's theory of sensory representation incompatible with Kriegel and Prinz's view. Kriegel and Prinz's view must be rejected, for the simple reason that its components are mutually inconsistent.

Now suppose that the proponent of the Kriegel-Prinz theory replies by arguing that in fact there is no incompatibility between a Tye or Dretske theory of representation and the claim that we sensorily represent response-dependent properties of the form *normally causing brain state B*. (This was suggested by Uriah Kriegel in discussion.)

For instance, the proponent of the Kriegel-Prinz theory might argue that (contrary to Tye's own view on the matter) Tye's theory of sensory representation is compatible with Kriegel and Prinz's claim that the represented properties are response-dependent properties, provided that we adopt an implausibly weak theory of "causally efficacy" on which the relevant "response-dependent" properties are causally efficacious, so that they can be represented.

Alternatively, the proponent of the Kriegel-Prinz theory who favors Dretske's theory of sensory representation might insist that our sensory systems have the function of indicating response-dependent properties of the form *normally causing brain state B*.

Now in my view these claims are totally incredible; I do not see how the proponent of the Kriegel-Prinz theory might defend them. For instance, what exactly is the relevant permissive theory of causal efficacy? Won't it be *too* permissive? I'm not sure I have much of a grip on the philosopher's notion of 'causal efficacy' but I thought it was supposed to be a *selective* relation between properties and effects. If this is denied I am not sure I have a grip on the notion at all. I find equally mystifying the claim that our sensory systems have the function of indicating response-dependent properties of the form *normally causing brain state B*.

But let's just pretend, for the sake of argument, that these incredible claims are true. Let's just pretend that there are suitable *permissive* versions of Tye and Dretske's psychosemantics. *Even if we are so concessive*, we can decisively rule out the Kriegel-Prinz theory. For now the proponent of the Kriegel-Prinz theory would face a different problem, *the profligacy problem*.

To see the problem, let's abstract way from the differences between the relevant permissive versions of Tye and Dretske's psychosemantics. Let *permissive tracking* be the relevant permissive relation, however exactly it is explained (in terms of permissive causation or permissive function). Now, if Yuck permissively tracks the property *normally causing A1*, then he presumably also permissively tracks the property *normally causing A2*, and the chemical property *C*, and so on. Likewise for Yum. So now we get the result that Yuck and Yum would both sensorily represent the *same* huge swarm of response-independent and response-dependent properties of the berry that they taste:

The chemical property C
Normally causing peripheral neural state P1 in humans
Normally causing peripheral neural state P2 in humans
Normally causing cortical neural state A1 in humans

Normally causing cortical neural state A2 in humans
Actually causing cortical neural state A1 in humans
Actually causing cortical neural state A2 in humans
Normally causing functional state F1 in some creature
Normally causing functional state F2 in some creatures

This absurd multiplication of represented properties is hard to accept. If the proponent of the Kriegel-Prinz theory says that Yuck and Yum sensorily represent all these properties, then I am not sure I understand what he might mean. More importantly: if Yuck and Yum represent the *same* huge swarm, it would leave us without an intentionalist account of the phenomenal *difference* between their taste experiences in terms of a representational *difference*.

At this point, the proponent of the permissive Kriegel-Prinz theory might say that Yuck and Yum *plain represent* the same swarm of properties by virtue of permissively tracking those properties; but they *sensorily represent* different properties; phenomenal character is determined by what properties are sensorily represented; so Yuck and Yum have phenomenally different taste experiences. (This was suggested by Uriah Kriegel in discussion.)

But then the proponent of the Kriegel-Prinz theory faces the *selection problem*. What is the reductive account of sensorily representing? What selects which of the (different) properties in the swarm Yuck and Yum "sensorily" represent? Of course, the selective sensorily representation relation cannot be identified with the permissive tracking relation. Maybe he could say that sensorily representing is some restriction on permissively tracking. But what is the restriction? Of course, he cannot say that X sensorily represents Y iff X permissively tracks Y *and* X's doing so determines the phenomenal character of his experience. For he is seeking a non-circular reductive account of phenomenal character (a similar point is made against Kriegel in Levine's NDPR review).

I conclude that the Kriegel-Prinz theory is mistaken. The main problem is the incompatibility problem, but even if we are concessive and set that problem aside, the theory still can be ruled out.

Let me enter a clarification. I'm *not* arguing against a response-dependent (or what Jonathan Cohen would call a "relational") view of the tastes of things or bodily pain-qualities. I'm only saying that one cannot combine this "response-dependent" theory with intentionalism and a Tye/Dretske theory of sensory representation, as Kriegel and Prinz attempt to do. As I have just argued, these views do not go together. However, for all I have said, one can combine the response-dependent ("relational") theory of sensible properties with other views on experience and sensory representation. In fact, I think one can, although I also believe that there are difficulties here and I do not myself endorse such a theory.[36]

[36] I have just said that I'm *not* arguing against a response-dependent (or what Jonathan Cohen would call a "relational") view of the tastes of things or bodily pain-qualities – I'm only arguing against the conjunction of this view with certain other claims. However, I do think that some who accept a response-dependent view face a psychosemantic problem, even if they reject these other claims. For instance, Cohen (2009) develops a response-dependent ("relational") account of color. And he says that our experiences represent fine-grained relational properties. But, as I point out in my (2010) review, he does not explain how. (Cf. also Hill 'Visual Qualia' against the view that shape experience represents relational shape properties – a view recently defended by Cohen which parallels his view of color experience.) In effect, he himself admits that he cannot use any of the usual

9. Concluding Remarks: A New Puzzle about Consciousness

I will conclude with some remarks concerning what the truth theory of conscious might look like, in the light of my arguments.

To begin with, let me say that I think that the arguments generalize. I think the arguments apply equally to conjunction of intentionalism about experience and other standard "externalist" theories of intentionality, for instance Fodor's asymmetric dependence theory. And, although I won't argue for this here, I think that the arguments I have developed here apply against the conjunction of intentionalism and a Millikan-style theory.[37] (To be clear: I don't think they make a problem for Millikan's theory of representation alone, but only for the conjunction of that theory of representation and the intentionalist view that all phenomenal facts are grounded in representational facts.)

What moral should we draw?

One very natural response would be to accept type-type identity (Churchlands, Clark, Block, McLaughlin, Hardin, Polger). This holds that experience properties like having pain-experience or having taste experience are necessarily identical with internal neural properties: S1 states, ensemble activations, and so on.. This theory accommodates internal-dependence. By stipulation, the individuals in coincidental variation cases, for instance Yuck and Yum and Mild and Severe, differ in their neural properties, even if they bear the same tracking relations to the environment. So, by contrast to tracking intentionalism, the identity theory easily accommodates the verdict that they have different experiences. The

---

psychosemantic theories to explain how our color experiences represent such properties. For concerning my Max-Twin Max case (a coincidental variation case involving color vision), he agrees with me (2009, 86) that those theories entail that they represent the *same (response-independent)* properties, whereas he wants to say (2009, 88) that they represent *different* fine-grained relational properties.

Here's a potential solution to the problem. Cohen accepts (2009, 214ff) a permissive, highly context-dependent theory of 'causal relevance' on which the truth-values of reports of causal relevance depend on our state of knowledge. On this view, even one of his relational properties R (having some properties that dispose its bearers to produce internal color experience state E) can be causally relevant to the production of our internal states (e. g. E). So you might think Cohen might combine this kind of permissive theory of causal relevance with a causal theory of how our experiences represent such relational properties.

But this combination of views is open to a *reductio ad absurdum*. Suppose you are looking at a tomato and that on this combination of views I can truly say that your experience represents R (because in this context R is causally relevant in the right way to the production of your experience, say). Now suppose my knowledge-state changes in such a way that I now cannot say that R is causally relevant to your experience. Now on this combination of views I would speak falsely if I said that your experience represents R. But there hasn't been an intrinsic change in *your* situation. This seems manifestly absurd. In short, if representation is explained in terms of causal relevance, and causal-relevance is highly context-dependent as Cohen says, then representation is highly context-dependent – which seems absurd.

Cohen actually faces this problem about *seeing*. Suppose you look at a tomato. Cohen explains what it is for you to see a relational color property of the tomato in terms of its being causally relevant (2009, 162, and note 14). So this view 'You see the redness of the tomato' is true in my context only if 'The redness [on Cohen's view a relational property] is causally relevant to your experience' is true in my context. Later (2009. 214ff) he says that whether I can truly say such a relational property is causally relevant depends on my epistemic state: there are special epistemic states in which I cannot truly say this. So, just because *my* epistemic state changes, 'You see the redness of the tomato' in my mouth might go from true to false, *even though your situation has not changed in anyway*.

[37] See Pautz section 8 of Pautz https://webspace.utexas.edu/arp424/www/simple2.pdf

identity theory might also accommodate our structure judgments, as we saw in connection with Max's judgment that his right arm pain is twice greater than his left arm pain and Harold's judgment that his .

So the identity theory easily accommodates Internal-Dependence. But I think we have to reject it because it does not accommodate what we might call External-Directedness. I have developed this argument elsewhere[38], so I will be brief. What External-Directedness comes to is that some experience properties are necessarily directed at shapes, locations, distances, spatial relations, and other properties which, if they are instantiated at all, are not instantiated in the brain. I think this applies as much to gustatory, olfactory and bodily experience properties as it does to visual experience properties. But neural properties do not have this modal property: they are not essentially external-directed in this way. So, while Internal-Dependence shows that the brain certainly plays a role in determining experience, we cannot accommodate this role by simply identifying experience properties with neural properties. A more nuanced view of the relations between the phenomenal and the neural is required.

In short: biological theories evidently accommodate internal-dependence. But they do not accommodate external-dependence. By contrast, standard, tracking forms of intentionalism err on the other side. They accommodate external-dependence, but they do not accommodate internal-dependence. So what we need is a middle way. We need a theory that accommodates both external-directedness and internal-dependence.

This leads to a puzzle, on that is virtually ignored. It is especially a puzzle for physicalists.[39] I think we should accommodate External-Dependence by accepting intentionalism: *all* sensory phenomenal facts are grounded in representational facts (with the possible exception of blur, etc.). So I think that if we wish to accommodate Internal-Dependence as well as External-Directedness, we must say that the individuals in coincidental variation case bear sensory representation relation to *different* external properties (phenomenal colors, felt pains, taste qualities, smell qualities) due to internal neural-functional differences, even though they bear the tracking relation, and other naturalistic relations, to the same external properties. These representational differences ground the phenomenal differences between them. The puzzle is: *what account of the sensible properties and the sensory representation relation might make it intelligible how this is so?* How can experience be a matter of presenting properties "out there" and yet what properties get presented depends on what happens inside the head. This is a puzzle for the simple reason that we have no ready model here.

But let me conclude by describing some possibilities.

Intentionalism plus Lewis-style functionalism-interpretationism: White and Jackson. It is part of the coincidental variation cases I described that the individuals involved have different behavioral dispositions on the output side, even if the track the same properties in the environment on the input side. Now *tracking intentionalists* cannot use this fact to argue that their experiences differ in representational content and hence phenomenal character. As we saw, tracking intentionalism entails that their experiences are exactly the same in representational content. But if we have a theory of intentionality that brings in output-factors in a more serious way, then we might think the behavioral differences between the

[38] See https://webspace.utexas.edu/arp424/www/bio.pdf

[39] As far as I know, only Colin McGinn alludes to the problem (in *Mental Content* someplace).

individuals involved ground intentional differences and thereby phenomenal differences. David Lewis had such a partially output-based theory. He was not committed to intentionalism. But Frank Jackson and Stephen White do accept intentionalism. And they apply Lewis's theory to sensory representation. Applied to sensory representation, the idea is roughly that you "sensorily represent" a property F iff the interpretation that does best at rationalizing your (and others') behavior assigns to your internal state an experience of an F. Since it is stipulated feature of coincidental variation cases that the individuals involved – Yuck and Yum, Harold and Henry, and Mild and Severe – differ in their behavioral dispositions on the output side, maybe the Lewisian "best interpretations" assign to their experiences contents involving different properties: e. g. different response-dependent (role) properties of the kind we looked at earlier in connection with the Kriegel-Prinz theory, or different "centering-features" as on Andy Egan's view (2010, 94), or different uninstantiated primitive sensible properties. So one might think this combination of intentionalism and Lewisian, partly output-oriented theory of sensory representation explains how sensory consciousness is "internally-dependent" as well as "externally-directed".

But Lewis's interpretationist-functionalist theory of intentionality faces well-known problems. I think that there are additional problems with the attempt to combine it with intentionalism about experience (cf. "A Simple View of Consciousness"). The problems are not the usual problems is psychosemantics. They are special problems that I think mean that this combination of views is not a serious option.

Fingers-crossed Reductive Intentionalism: Lycan. William Lycan accepts intentionalism and is inclined to accept what I've called Internal-Dependence. That is, he thinks that in some coincidental variation cases the right verdict is Different Experiences and hence, given intentionalism, Different Contents. So he is inclined to accept the argument I have given against tracking intentionalism.[40] But he does not attempt to describe how a reductive version of intentionalism might deliver the verdict of Different Contents and hence Different Experiences in coincidental variation cases.[41] He admits that tracking intentionalism won't. In fact he thinks no one has come up with a decent reductive version of intentionalism about experience. Yet he is still a reductive intentionalist; so he has what we might call a 'finger-crossed' reductive intentionalist.

[40] See Lycan https://webspace.utexas.edu/arp424/www/vs.pdf Lycan writes: "All [Pautz] needs [to undermine Dretske and Tye's tracking intentionalism] is the judgment that [in coincidental variation cases] Different Experiences is by far more plausible than Same Experience, and with the latter claim I tend to agree". This seems to require that he reject his earlier view (1996, 75) that phenomenal character is determined by the representation of response-independent properties of external things, for on any reasonable view the individuals in coincidental variation cases represent the same such properties.

[41] Lycan (ibid) says Millikan's theory solves some problems with tracking theories, and says that it follows that it is reasonable to think that the intentionalist might be able to use it to provide an account of coincidental variation cases (although Lycan himself thinks that all extant psychosemantic theories, including Millikan's, are ultimately unsatisfactory). But, I don't see how this follows. (Even if solves some problems, why think it solves others?) And, he doesn't explain how the intentionalist might use it to provide an account of coincidental variation cases. Finally, I elsewhere argue that Millikan's theory is of no help: section 8 of https://webspace.utexas.edu/arp424/www/simple2.pdf

Is this attitude reasonable? Of course, this question is beyond the scope of this paper. But let me just register one point that makes it seem that this is not a reasonable stance – elsewhere I develop some more systematic considerations.[42]

To illustrate the point, consider a coincidental variation case in the abstract. We have two individuals, A and B, who generally have different experiences through time. A has E1, E2, E3, etc. and B has E1*, E2*, E3*, etc. Yet they optimally track the same response-*independent* properties. Now if we are reductive intentionalists we need to do two things: we need to specify the properties they sensorily represent, and we need to give a reductive account of how they sensorily represent different such properties (thereby having different experiences). But what are those different properties? Not the response-*independent* properties, obviously. Not properties like *being bad:* such properties are not fine-grained enough to determine the relevant phenomenal differences and in any case this idea does not apply to all coincidental variation cases, as we have seen (see my reply to the fifth objection in §7). If we are reductionists, then the only candidates appear to be Shoemakerian response-dependent properties: normally causing internal state B1, normally causing B2, normally causing B3, and so on. These response-dependent properties are fine-grained; the representation of them might determine all the fine-grained phenomenal differences. If we hold that A and B have different experiences, we accept intentionalism (all phenomenal differences grounded in representational differences), and we want to be reductionists, then this response-dependent gambit seems to be the only way to go. Lycan (2006, 4.4.3) himself has taken seriously a similar Shoemaker-style response to Levine's fish creature argument (although, as I have said (§2 and footnote 8), I think that he might also just respond to *that* argument by saying that the scenario is not possible).

But what *reductive* theory of the sensory representation relation might explain how we represent such funny properties? Shoemaker has no such theory.[43] Kriegel and Prinz do, but it does not work, as we have seen. Their theory faces the incompatibility problem, the profligacy problem, and the selection problem. And I do not see how these problems go away if we consider other reductive theories of intentionality. The problem here is not just one of the usual problems in psychosemantics: the problem is that all psychosemantic theories that we know about are clearly incompatible with the claim that we *only* sensorily (phenomenally) represent some sort of such peculiar response-dependent properties (and not other such properties).

So if we accept intentionalism and internal-dependence, then we have reason not to take the fingers-crossed attitude that there is some reductive version of intentionalism out there that solves our problems. But I definitely think it is worthwhile to attempt to find one.

Non-reductive Intentionalism (David Chalmers, Mark Johnston, Colin McGinn). On intentionalism, episodes of sensory consciousness involve a dyadic relation between individuals and complexes of properties apparently present in external objects or bodily regions: the sensory representation relation. We might also call it the "presentation relation". According to non-reductive intentionalism, we ought to give up the goal of reducing this relation to a physical-functional

[42] https://webspace.utexas.edu/arp424/www/simple2.pdf

[43] For more on why Shoemaker does not have a reductive theory of the sensory representation relation, see Pautz section 5 of https://webspace.utexas.edu/arp424/www/rest.pdf

relation. Same for sensible properties. Yet we can accommodate internal-dependence by claiming that the "presentation relation" supervenes on internal-functional factors. (Compare Horwich *Meaning*; Shoemaker, in discussion.) I believe that, *if* non-reductive intentionalism is true, it has important consequences concerning the mind-body problem and our understanding of non-sensory forms of intentionality.

To sum up. The usual arguments against tracking intentionalism based on intuitions about far out cases – inverted spectrum, inverted earth, brain in vats, and so on. Tracking intentionalists reject the arguments by rejecting the intuitions they're based on. I've attempted to break the impasse by using arguments that are firmly rooted in empirical research. Kicking qualia "downstairs" flies in the face of decades of research in psychophysics and neuroscience. One response would be to accept type-type identity theory and reject intentionalism. The response I favor is to keep intentionalism but try to come up with a more "internalist" theory of sensory intentionality. On this view, sensory consciousness is externally-directed as well as internally-dependent. But it is puzzling how this might be so.

William Lycan (an honest philosopher) has said "materialist psychosemantics is in very bad shape and not getting better" (2008). I tend to agree with that assessment. But I would also like to make a plea for at least giving it a shot: for attempting to really work out in detail a reductive version of intentionalism across the sense-modalities that is in line empirical facts (in addition to worry about more high-level issues like Zombies, the Knowledge Argument, what distinguishes conscious states in general from non-conscious ones, and so on): something that has not really been done so far, in my opinion. Otherwise why should we put any stock in this view?[44]

References

Armstrong and Malcolm, 1984. Consciousness and Causality
Block 2005 An Obstacle to Representationalism about Pain. In Aydede.
Byrne 2009 Experience and Content, PPQ
Byrne 2003 Color Similarity, PPR
Byrne and Hilbert 2003. Color Realism and Color Science BBS.
Byrne and Hilbert 1997 Colors and Reflectances
Byrne and Tye 2006 Qualia Ain't in the Head
Coghill et al. (1999) Pain Intensity Processing Within the Human Brain: A Bilateral, Distributed Mechanism
Coghill 2003 Neural correlates of inter-individual differences in the subjective experience of pain.
Cohen 2009 The Red and The Real
Cowart and Rawson (2005) Article on smell in Goldstein, Handbook on Sensation and Perception
Dretske 2003 How Do you Know you are not a Zombie?
Gescheider 1997 Psychophysics: the Fundamentals
Harman 1996 Explaining Objective Color in terms of Subjective Reactions, *Phil Issues*.
Hofbauer. Cortical Representation of the Sensory Dimension of Pain
Howard and coworkers on *human subjects* (2009)xx

[44] Hill's *Consciousness* attempts to provide a reductive theory of intentionalism about vision and pain which actually takes into account some of the empirical complexities – although as I have explained here and elsewhere (https://webspace.utexas.edu/arp424/www/hill.pdf) my opinion is that, despite Hill's impressive efforts, in the end his reductive-externalist intentionalism does not work. O'Challaghan and Batty have given us very interesting discussions of sound and smell but again I think that the sorts of considerations I develop here might raise problems for their views.

Hill 2005 Ouch! In Aydede.
Hill 2006 Perceptual Experience: How it Opens the World. In Kriegel *Self Representational Approaches to Consciousness*
Jackson 2004 Experience and Representation
Johnston, Mark. Objective Minds. PPR
Kenshalo 2000 Response Properties and Organization of Nociceptive Neurons in
Area 1 of Monkey Primary Somatosensory Cortex
Levine, J. (1988). "Absent and Inverted Qualia Revisited", Mind & Language, vol. 3, no. 4, 271-287.
Levine, J. (2003). "Experience and Representation", in Consciousness: New Essays, Quentin Smith and Aleksandar Jokic eds., Oxford University Press.
Linster and coworkers (2001) Perceptual Correlates of Neural Representations Evoked by
Odorant Enantiomers
Lycan 2006 Representational Theories of Consciousness SEP.
Lycan 2008 Phenomenal Intentionalities.
Malnic et al. 1999) Combinatorial Receptor Codes for Odors
Margot (2009A noseful of objects
Pautz 2010 Review Cohen The Red and The Real Notre Dame
Pautz 2006 Can Color Structure Be Explained in terms of Color Experience? AJP.
Pautz forthcoming a. Can Disjunctivists Explain Our Access to the Sensible World? Phil Issues 21
Pautz forthcoming b. Phenomenal Intentionality: A Backhanded Defense. Forthcoming in Phenomenal Intentionality (Kriegel and Horgan) .
Porro 1998, Temporal and Intensity Coding of Pain in Human Cortex
Price, D. 1999. *Psychological Mechanisms of Pain and Analgesia:* Seattle: IASP Press.
Price 2002 Central Neural Mechanism that Interrelate Sensory and Affective Dimensions of Pain
Prinz 2010 When is Perception Consciousness? In Nanay Perceiving the World
Prinz 2006 Churchland on State Space Semantics
Prinz 2006 Beyond Appearances
Speaks, Jeff. A quick argument against phenomenism, Fregeanism,
appearance property-ism, and (maybe) functionalism about perceptual content.
Smith, D., R. Van Buskirk, J. Travers, and S. Bieber. 1983. Coding of taste stimuli by hamster brain stem neurons. *Journal of Neurophysiology* 50: 541-558.
Stich, 1990 PPR discussion of Dretske *Explaining Behavior.* Building Belief. *Philosophy and Phenomenological Research* 50: 801-6.
Stevens et. al (1958) A SCALE OF APPARENT INTENSITY OF ELECTRIC SHOCK
Timmerman et al. (2001)Differential Coding of Pain Intensity in the Human Primary and Secondary Somatosensory Cortex
Tye 1995 Ten Problems
Tye 2000 Color Consciousness and Content
Youngentob, Predicting Odorant Quality Perceptions from Multidimensional Scaling of Olfactory Bulb Glomerular Activity Patterns