

The Real Trouble for Phenomenal Externalists:

New Empirical Evidence for a Brain-Based Theory of Consciousness

Adam Pautz

[We should] reverse the whole programme started by Galileo - *we should put these [sensible] qualities back into the physical world again.*

--David Armstrong

The traditional view of the sensible qualities locates them in the head. But within philosophy there has recently been a kind of externalist revolution. While most scientists would still locate the sensible qualities in the head, many philosophers now claim that sensible qualities are really “out there” in the mind-independent physical world and that the function of the brain is just to reveal them to us. In favorable conditions sensory character is determined simply by what mind-independent states you are directly conscious of. The result is a kind of phenomenal externalism. Examples include externalist intentionalism, naïve realism, and active externalism.¹ The stakes are high, because many think that phenomenal externalism represents our best shot at naturalizing consciousness and its intentionality.

My own view is that phenomenal externalism has been a big wrong turn. I favor a kind of internalist counter-revolution. But, for reasons I will explain, I disagree with those who think phenomenal externalism can be refuted very easily on the basis of controversial intuitions about brains in vats (Horgan, Tienson and Graham), inverted spectrum (Shoemaker), actual cases of perceptual variation (Block), and so on.²

¹ For active externalism, see Noë (2004) and O'Regan (2011). Although these authors advertise active externalism as a radical kind of phenomenal externalism (phenomenology fails to supervene on the brain), this is sometimes unclear. For instance, in response to an objection from David Chalmers, Noë (2004, p. 119) changes his view, claiming that his view is that phenomenology is constituted by a slew of sophisticated *beliefs* or *expectations* concerning what the sensory effects of various actions would be. In that case, his view might actually be a version of phenomenal internalism, because (for all he says) the relevant beliefs might be narrow beliefs that supervene on the head. In general, as Block (2012) shows, active externalists do not have any clear view. Partly for this reason, here I will be focusing on other varieties of phenomenal externalism.

² See Horgan, Graham and Tienson (2004), Shoemaker (1994) and Block (1999).

Both sides of the debate have missed the best argument against phenomenal externalism. The real trouble with phenomenal externalism is that it goes against decades of research in psychophysics and neuroscience. The basic point is that, even under ideal conditions (no interfering factors), sensory character is much better correlated with neural patterns in the brain than with anything in the external, physical world. I call this the *problem of correlations* for phenomenal externalism. For this reason most scientists of sensation and perception would probably not take very seriously the kind of phenomenal externalism now being promoted by some philosophers. I have begun developing the argument in previous work. In the present essay, I will go beyond that work, by using impressive new research and by ruling out recent responses.³

To make the discussion concrete, I will initially focus on what I call “tracking intentionalism” as a kind of stalking horse. This is a version of externalist intentionalism, which combines the intentionalist thesis that all phenomenal differences among sensory experiences are representational differences with a reductive externalist theory of representation. But while I will focus on tracking intentionalism I would like to stress that I believe my arguments will undermine any reductive variety of externalist intentionalism. Defenders of views in the general vicinity include David Armstrong, Alex Byrne, Fred Dretske, David Hilbert, Christopher Hill, William Lycan, David Papineau, and Michael Tye. I agree with these philosophers that sensory experiences are intentional states that present sensible qualities ostensibly located in the external world or bodily regions; what I will argue against is only their *externalist* variety of the view. Although I will not discuss this here, I believe my empirical arguments also undermine the version of phenomena externalism defended by John Campbell and other “naïve realists”.⁴

Many externalists focus narrowly on one sense-modality, without showing how externalism can be developed across the board. To show just how unpromising the externalist picture is, I will consider multiple sense-modalities. In particular, I will target recent externalist views of

³ See Pautz (2006a) and (2010). Hill (2012) provides very interesting externalist responses to the kind research on pain mentioned in Pautz (2006, p. 212-213) and elaborated in Pautz (2010) and in the present paper. Cutter and Tye (2011) also reply to my empirical objections about pain. I will take their view into account throughout this paper.

⁴ For externalist intentionalism, see Armstrong (1999), Byrne and Hilbert (2003), Dretske (1995), Hill (2009), Lycan (2001), Papineau (2012), and Tye (2000). For naïve realism, see Campbell (2002).

taste qualities (Smith), smell qualities (Batty), auditory qualities (O'Callaghan), and pain qualities (Tye, Dretske, Hill).

My plan is as follows. In §1 I explain tracking intentionalism. In §2 I show that, even under ideal conditions, sensory character is much better correlated with neural patterns in the brain than with anything in the external, physical world. In §§3-4 I elaborate in detail two independent empirical arguments against tracking intentionalism based on the correlational data. In §5 I eliminate recent responses which involving defending more elaborate versions of externalist intentionalism. These responses appeal to functional or syntactic features of representations (Lycan, Hill), binding-errors (Hill), valuational or threat-level contents (Cutter and Tye, Hill), complex properties (O'Callaghan), or response-dependent properties (Kriegel). Finally in §6 I will suggest that, although *externalist* intentionalism fails, sensory experiences are indeed nothing but intentional states that present sensible qualities ostensibly located in the external world or bodily regions. However, for empirical reasons, the best way of developing intentionalism is by accepting what David Chalmers called an “Edenic theory” of sensible qualities.

1 What Is Tracking Intentionalism?

I start by explaining the basic tracking intentionalist picture that will be my stalking horse. It has three parts.

The *first part* is a reductive and objectivist theory of sensible qualities like colors, smell qualities, taste qualities, audible qualities, and so on. It is *reductionist* in that it holds that sensible qualities are physical properties, in a suitably broad sense of ‘physical properties’. It is *response-independent* in that it holds that sensible qualities are not in any way to be defined in terms of effects on perceivers.

So, for instance, maybe colors are reflectance properties, smell qualities and taste qualities are chemical properties of odor clouds and foods, auditory qualities are extremely complex physical properties involving frequency, amplitude, duration, and “critical bands”. And maybe perceived shapes are viewpoint-relative but objective properties like *being-elliptical-from-viewpoint-p* and *being-round-from-viewpoint-q* (Hill 2009). Now it is well known that even under optimal conditions multiple physical properties can cause one to be ostensibly conscious of the same sensible quality. In color vision this is known as “metamerism”. The same phenomenon occurs in all the sense-modalities. But

tracking intentionalists are not perturbed: they simply reduce a sensible quality to the disjunction of all the physical properties that normally give rise to our experience of it.

Tracking intentionalists even say that pain qualities you feel in your body are mind-independent physical properties, for instance, types of bodily damage or “potential” damage. This view faces what I have called the “*percipi puzzle*” (Pautz 2010). Byrne (forthcoming) has more recently called it the “*puzzle of pain*” in his discussion of Hill’s (2009) somewhat different paradox of pain. On tracking intentionalism, just as colors can exist without experiencers, so can felt *pain qualities*! For pains as well as colors are treated as entirely mind-independent physical properties of external items. For instance, the *very same horrible quality* you feel in your thumb when you hit yourself with a hammer (on, this view, a kind of damage) might occur in an insentient cadaver! However, here I propose to set *a priori* objections to the side. My aim will be to develop new empirical problems.

The *second part* of tracking intentionalism is a broadly *tracking theory* of sensory awareness of *properties*. The rough idea is that you sensorily represent an objective sensible quality (on this view, a physical property), and are thereby aware of it, just in case you undergo an internal state (a “*representation*”) that “*registers*” or “*tracks*” the instantiation of that property by external items. By using the term ‘tracking’ I do not mean to presuppose a simple input-based, causal theory of representation. I use “*tracking*” in a totally neutral way, as a kind of placeholder for a more detailed story.

However, for the purposes of illustration, I will largely focus on views of Michael Tye and Fred Dretske. Tye (2000) reduces the sensory representation relation to the *optimal tracking relation*, that is, the relation: individual *x* is in an internal state that plays 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*). Dretske (1995) 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*. By ‘functional role *F*’ I mean the special functional role that is supposed to turn unconscious representational states into consciousness ones: maybe some kind of cognitive accessibility. This will not concern us here. While I focus on Tye and Dretske, we will see (in §5) that my arguments also work against more complex versions of externalist intentionalism, including those which

appeal to Millikan's (1989) consumer-based approach to representation.

The *third part* of tracking intentionalism is intentionalism about sensory phenomenology. At a minimum, the intentionalist says that if two individuals are ostensibly conscious of, or sensorily represent, the very same sensible qualities (at the same places), then they have phenomenally identical sensory experiences. In cases of illusion and hallucination, the presented sensible qualities do not belong to any external items. Sensory content determines sensory phenomenology. This, or something like it, is extremely plausible. For instance, intuitively, if two individuals are ostensibly conscious of the very same smell or taste qualities, then they must have phenomenally identically smell or taste experiences. If they are ostensibly conscious of the very same auditory properties (from the apparent same direction), they must have phenomenally identical auditory experiences. Intentionalism is just a theoretical gloss on these intuitions.

That, then, is tracking intentionalism. It is undeniably attractive. Indeed, Cutter and Tye (2011, p. 91) have recently said, "tracking [intentionalism] is the most promising view for the philosopher in search a naturalistic account of experience". The reason is simple. Sensory consciousness is *externally-directed*. The sensible qualities certainly *appear* to be out there, in objects, in bodily regions, in foods (or maybe the tongue), and so on. This favors "objectivism" about sensible qualities and sits poorly with a Galilean view that locates the sensible qualities in the head. And if objectivism is true, then tracking intentionalism, or something like it, appears almost inevitable. For, in order to explain in naturalistic terms how the mind can become aware of, or "represent", objective properties out there, objectivists must presumably appeal to *causal* or *indicator* or *teleological* relations between brains and those properties. What other option is there? What else could hook us up to these properties? In that case, what sensible qualities you perceive are fixed by extrinsic factors, namely, your relations to your environment. Since sensory phenomenology is intuitively inseparable from what sensible qualities you perceive, the result is a radically externalist theory of phenomenology. So, while objectivism is a theory of the sensible qualities and tracking intentionalism is a theory of phenomenal character, I think that anyone attracted to objectivism about sensible qualities is under pressure to accept an externalist theory in the vicinity of tracking intentionalism, including for instance Casey O'Callaghan (2002) and Clare Batty (2010).

Of course, the theory I have presented is very simple. Many would like to add some bells and whistles. But I think many are committed to externalist views of sensory character in the general vicinity, including David Armstrong, Alex Byrne, Fred Dretske, David Hilbert, Christopher Hill, William Lycan, David Papineau, and Michael Tye.

Let me mention two caveats. First, Lycan (forthcoming) and Hill (2011) seem to hold that some phenomenal differences among sensory experiences are grounded in functional-syntactic differences, not representational differences. But we will see (in §5) that this cannot save their views from my arguments. Second, Byrne and Hilbert (2003) as well as Hill (2009) express skepticism concerning all existing naturalistic theories of representation. But they still hold that sensory qualities are physical properties of external things, and that phenomenal character is (at least largely) determined by the representation of these properties, and that some externalist naturalistic theory of representation is correct (even if we cannot specify it). As we shall see, this is enough to make them vulnerable to my arguments.

So scores of philosophers take the same basic externalist approach. And it is very attractive, because it fits with the externally-directed character of sensory consciousness. The only trouble is that it flies in the face of decades of research in psychophysics and neuroscience. That research shows that consciousness is (in a non-trivial sense) internally-dependent, even if it is also apparently externally-directed.

2. It's Only in Your Head: The Neural Basis of Some Phenomenal Facts

Tracking intentionalism is radically externalist. On tracking intentionalism, the character of experiences is not determined by the intrinsic character of their neural correlates. Here is an analogy. The shapes of words do not matter to what they represent. Thus, 'dog' and 'cat' are physically dissimilar, but represent similar animals. In general, there is a sense in which the intrinsic features of the neural content-carriers do not matter to what contents they carry. Likewise, on tracking intentionalism, there is *some* sense (which I will make more precise in §4) in which the intrinsic features of postreceptor neural processing do not matter to phenomenal character. All that matters to phenomenal character are what physical properties the neural wetware tracks and thereby represents.

But this could not be more wrong. In fact, the whole history of psychophysics and neuroscience shows that *exactly the opposite is true*. In *some* cases, the intrinsic features of the neural wetware does somehow matter, in a way that I will later show to inconsistent with tracking intentionalism (§3). In particular, two facts are relevant.

First, for decades psychophysics has revealed that, even under optimal conditions (no interfering factors), there is *some* sense in which there is an extremely bad correlation between experiences and external physical properties tracked. What I mean will become clear as when consider examples. But let me say at the outset that I do not merely mean that the physical properties tracked are disjunctive or unnatural (because of metamerism). Roughly, what I mean is that, even under optimal conditions, the *structural relations* among experiences (similarity and difference, equal intervals, proportion) are not matched by the *structural relations* among the (disjunctive) external physical properties that those experiences track. True, in some cases, they do match; in other words, there is good external correlation. For instance, under optimal conditions, subjects' reports on when perceived length doubles corresponds to an actual doubling in physical length. But psychophysics has shown that this is the exception rather than the rule. When it comes to taste, smell, pain and sound, there is *bad external correlation*. Here tracking intentionalists have it exactly wrong. The external physical world is just the *wrong place* to look for the basis of qualitative character.

Second, neuroscience has revealed that experiences are much better correlated with neural firing patterns in the brain. What I mean by this, too, will become clear as we go on. But let me say at the outset that I do not merely mean that every distinct experience is co-extensive with a distinct neural correlate in humans, so that every for measurable change in experience there is some measurable change in the nervous system. Some philosophers think that this is enough to refute tracking intentionalism. This is a mistake. It is equally true that going from *thinking about water* to *thinking about aluminum* requires a neural change; but no one would think this undermines externalism about thought about natural kinds. What I mean by good internal correlation is something subtler than the existence of correlations between individual experiences and individual neural states. What I mean is that in *some* cases *structural relations among* experiences (similarity and difference, equal intervals, proportion) are well matched by *structural relations among* their neural correlates. In these cases, while

there is bad external correlation, there is *good internal correlation*. In these cases the basis of certain structural facts about phenomenal character are to be found *only in the brain*.

To illustrate these points, I will in the rest of this section provide some empirical background concerning the fascinating science of taste, smell, pain and sound. I will wait until later sections to explain how the science can be definitively shown to be at odds with tracking intentionalism.

2.1 Taste

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. All sensory-phenomenal facts about taste experiences are determined by the chemical types they optimally track and thereby represent.

The trouble is that phenomenal *resemblances and differences* among taste experiences are not well correlated with resemblances and differences among the chemical types they track. Examples of bad external correlation abound. For instance, suppose you taste aspartame and then a stereoisomer of aspartame. The chemical properties that your taste experiences optimally track are *extremely similar*: the compounds only slightly differ in the orientation of two hydrogen atoms. Yet your taste experiences are *extremely different*: the taste of aspartame is sweet while the taste of the stereoisomer is extremely bitter (Walters 1996). Likewise, gentiobiose is bitter, while trehalose has a distinctly sweet taste, even though they are very similar disaccharides composed of two glucose units. Indeed, gentiobiose has an *anomer* (a kind of very similar stereoisomer), namely isomaltose, which tastes sweet (Sakurai 2010). Neohesperidin, which is found in citrus peel, is extremely bitter; removing a single carbon-oxygen bond produces neohesperidin dihydrochalcone, which is extremely sweet.

These are examples of phenomenal difference despite chemical similarity. There are just as many examples of phenomenal similarity despite radical chemical difference. 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 expe-

riences. This makes evolutionary 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, chemical similarities and differences just cannot explain phenomenal similarities and differences. What then explains them? Neuroscience has shown that there is good “internal” correlation: in many cases, the phenomenal similarities and differences in taste experiences are better correlated with similarities and differences among neural states in the taste system than they are with anything in the physical world.

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, umami). But at more central locations in the taste system neurons are *broadly turned*, 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 *ensemble activation*.

The resemblance ordering among tastes is well correlated with the resemblance ordering among such ensemble activations (as determined by multidimensional scaling) in “neural similarity space”. There are a number of studies that bear this out. 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 general, the greater the phenomenal similarity, the greater the ensemble activation similarity. Likewise, Smith and coworkers (1983) created a multidimensional *neural similarity space* of ensemble activations in the hamster taste system in which distances among points represent degree of similarity. They found that 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.

These studies suggest that *spatial pattern* codes for taste quality: taste is coded by which neurons are activated and to what degree. But more recent studies show that the *temporal pattern* of firing within single neurons also contributes to taste coding. For instance, Di Lorenzo *et al.* (2009) found that the distinctive temporal patterns in nucleus of

the solitary tract (NTS) corresponding to *basic tastes* are very dissimilar, in a way that mirrors those tastes' phenomenal dissimilarity. Further, binary mixtures produce temporal patterns that are *in between* those produced by their respective components. The temporal patterns produced by mixtures were even typically well approximated by a linear superposition of those produced by their components. Indeed, Di Lorenzo *et al.* report that "the entire [three-dimensional] taste space can be mapped by the temporal characteristics of response in a single cell" (9232).

Of course, taste quality might be coded by both spatial pattern (ensemble activation) and temporal pattern. As Chen *et al.* (2011) write, "the existence of consistent temporal profiles of response among the responsive neurons for a given taste stimulus enhances the uniqueness of the across-neuron pattern of response by adding a dynamic dimension . . . thus the spatial pattern produced by a tastant is sculpted as the response unfolds over time".

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 and phenomenal data* (see also Oakley 1985).

2.2 Smell

Now let us turn to smell. 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 chemical type it optimally tracks.

But, in the case of smell, examples of bad correlation between experiences and the physical properties optimally tracked are even more plentiful than in the case of taste. Cowart and Rawson (2005, p. 568) 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. 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.

Yet on tracking intentionalism the chemical properties represented by our experiences are supposed to be what fully determines all aspects of the experiences' phenomenal character, including the phenomenal resemblances and differences between them.

Another interesting psychophysical fact is worth mentioning. The tracking intentionalist would presumably say that the *phenomenal intensity* of a smell experience is constituted by the particular *concentration* of the external odorant that it represents. This fits many cases, because in general changes in the level of concentration go with changes in stimulus intensity. But there are counterexamples. In some cases a mere change in the concentration of a chemical can strikingly alter the *quality* and not just *intensity* 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 constitute changes in intensity, sometimes changes in quality? In the external world, there is only a difference in degree; but in some cases the quality changes in a categorical way. As we shall see, neuroscience provides the answer. The puzzle is resolved by a more internalist view on which internal neural factors play a role in determining phenomenal character, in a way at odds with tracking intentionalism.

If bad external correlation means that the chemical properties optimally tracked by our gustatory experiences don't explain the phenomenal character of those experiences, then what does explain it? Neuroscience has revealed "good internal correlation", suggesting the explanation is to be found in the brain.

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

As noted, chemicals that are nearly identical structurally can elicit very different smell experiences. 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.

Indeed, even some enantiomers (chemicals that are mirror image rotations) can smell quite different to us, while others smell the same. For instance, –carvone smells like spearmint while its mirror image +carvone smells like caraway. The best 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) also produce quite different neural patterns further downstream in the olfactory bulb of *rats*. And those which smell the same produce similar patterns.

There is a striking demonstration of this kind of “good internal correlation” in the case of smell 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 Figure 1).

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). In particular, Howard and coworkers found that, even though the molecular structures in each of the three families are quite different, they produce very similar ensemble activations in PPC, which are distinct from the activation patterns of the other two categories (Figure 1). They even located ensemble patterns in a *three-dimensional neural similarity space*, and found that neural similarity (represented by distance) coincided very well with phenomenal similarity.




Figure 1. Howard *et al.* (2008) found that molecular structures that are chemically very different but smell similar (citrus, minty, woody) produced very similar ensemble activation patterns in the PPC. In general, degree of phenomenal similarity or difference coincided with similarity or difference in PPC patterns, not similarity/difference in molecular structure. Reprinted from Margot (2009) with permission.

Now I can explain why 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 their ensemble activation representations in the brain are not mere artifacts of similarities in the molecular structures of the odorants used (since they were *not* similar); the only explanation of the observed correlation be-

tween 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 a discussion of the Howard experiment:

[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]*. (My italics.)

The Howard study on humans is not the only study showing that similarity and difference in olfactory experience to be correlation similarity and difference in ensemble activations. There are also many studies on animals that show this as well. For instance, Youngentob and coworkers (2006) did a similar study on rats. Of course, by contrast to humans, rats unfortunately cannot report on the degrees of similarities among their smell experiences of chemicals. So to get at the phenomenal structure of their smell experience, a more indirect method is required. Youngentob and coworkers had the rats perform a *confusion matrix task*. The basic idea is that degree of phenomenal similarity corresponds to probability of confusion. Then, using a 2-DG functional mapping technique and multidimensional scaling, they looked at the degrees of neural similarities among the neural ensemble activations set up by the odorants in the rats' olfactory bulb. Here is what they found:

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 [multidimensional scaling] 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 relationships 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. . . (p. 1343; my italics)

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.

2.3 Pain

Suppose you have a variety of pains of different intensities in different 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 phenomenal aspect of the pain reduces to some physical feature of the bodily disturbance represented by the pains. So felt location is just represented location; and differences in *quality* among pains (prickling, stabbing, throbbing, etc.) are constituted by differences in the *types* of bodily disturbance they represent. Now, besides quality and location, you can also focus on the *sensory intensity* of a pain. (This related to but distinct from the *unpleasantness* of the pain, or the *affective dimension* of pain, as we shall see.) On tracking intentionalism, what (possibly complex) aspect or feature of the external stimulus constitutively determines its sensory intensity? As far as I know, tracking intentionalists have simply not addressed this issue. The simplest view would be that the sensory intensity of a pain is fully determined by the *intensity* and *size* of bodily disturbance optimally tracked and so 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. Stevens *et al* (1958) showed this in the case of electric shock. Likewise heat-induced pain intensity is a power function of stimulus-temperature, with an exponent of about 3.2 (see Figure 2). A small increase in tem-

perature can double subjects' ratings of the sensory intensity of pain. Pain intensity is not only a function of stimulus intensity; it is also a function of stimulus area, in a way that cannot be easily summarized (Price 1999).




Figure 2. The relationship between noxious temperatures and pain sensation intensity ratings is described by a power function with an exponent of about 3.2. Reprinted from Price et al. (1994).

So psychophysics has revealed that, even under optimal condition, there is a messy relationship between pain intensity and the many aspects of the bodily disturbance tracked. By contrast, many researchers have reported a *perfect correlation* between pain intensity and a single neural parameter, namely firing rates of neurons in the areas of the brain involved in pain. For instance, using noxious 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). (p. 1936)

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*. 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 the firing rates of wide-dynamic range (WDR) neurons in monkey S1 very closely resembles the psychophysically-derived relationship between temperature and pain intensity in humans shown in Figure 2. Indeed, just as pain intensity doubles between 47°C and 50°C, so WDR response in monkeys roughly doubles between these temperatures (see Figure 3).




Figure 3. The relationship between temperature and the firing rates wide-dynamic range (WDR) neurons in monkey S1 very closely resembles the psychophysically-derived relationship between temperature and pain sensation intensity in humans (shown in Figure 2). Reprinted from Kenshalo *et al.* (2000) with permission.

These experiments and a slew of other empirical considerations (Price 2002) suggest that S1 plays a special role in determining pain intensity. For the sake of simplicity, I will sometimes assume this view in what follows; but my arguments would go through on a more distributed view of the neural basis of the sensory intensity of pain as well (e. g. Coghill *et al.* 1999).

So far I have discussed the sensory dimension of pain. Pain of course also has an affective-motivational dimension. Price (2001, 393) describes it as “the moment-by-moment unpleasantness of pain, which consists of emotional feelings that pertain to the present or short-term future, such as annoyance, fear, or distress”. Many studies (e. g. Rainville *et al.* 1997; Hofbauer *et al.* 2001) actually show that sensory intensity and pain affect can be modulated independently and that, while sensory intensity is based in S1, the affective dimension of pain is based in the anterior cingulate nucleus (ACC).

So, in the case of pain intensity, while there is bad external correlation, there is very good internal correlation. There is messy, uncodifiable relationship between the size and intensity (e. g. temperature) 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 at the sensory level.

2.4 Audition

Finally, there is evidence of “bad external correlation” and “good internal correlation” in the domain of auditory perception.

Here are some examples of bad external correlation in the auditory domain (Moore 2003). There is a relationship between perceived *intensity* of a sound and the amplitude of the sound. But it is one of response *compression*. So, for instance, doubling perceived intensity requires *far* more than doubling amplitude. However, at lower amplitudes, loudness increases more rapidly with increasing amplitude. Perceived loudness is a function not just of amplitude but also of frequency, in a way that resists codification. For complex tones, loudness also depends on bandwidths (as we shall see, this has a cortical explanation). Similarly, there is a relationship between the perceived *pitch* of a sound and the frequency (for complex sounds, fundamental frequency) of the sound. But it too is one of response compression, with the degree of compression depending on frequency level. Perceived pitch also depends on amplitude. In particular, the pitch of tones below 2 kHz increases with increasing amplitude and the pitch of tones above 4 kHz decreases with increasing amplitude. For complex tones, pitch depends on a variety of other factors.

While there is bad external correlation in the auditory domain, there is some evidence of good internal correlation. Let me focus on the case of loudness, because the neural basis of pitch perception remains relatively poorly understood. Relkin and Doucet (1997, 2738) write that “the perceived loudness of a pure tone appears to be linked both to the number of spikes fired by single neurons and to spatial spread of excitation in the auditory nerve”. Langers et al. (2007) used fMRI to look at neural activity further downstream in the auditory cortex. They found that “cortical activity is more closely related to the perceptual loudness level of sound than to its [external, physical] intensity level” (p. 714) and indeed report “a type of non-linearity . . . comparable to that reported in psychophysical studies on loudness perception that employ subjective loudness scaling” (p. 716). On the

basis of this study and others, Röhl *et al.* (2011, p. 1494) conclude that “the most simple interpretation would be, that AC [auditory cortex] is fed by . . . the auditory brainstem according to the sound pressure level and the bandwidth of the stimuli, and an additional component is added which is linearly related to the perceived loudness”.

So the situation regarding loudness may be similar to the situation regarding pain intensity. In both cases, there is a non-linear relationship between sensory intensity and the physical stimulus. The difference is that, while in the case of pain intensity the relationship is one of *response expansion*, in the case of sound intensity it is one of *response compression*. The explanation is that there is also a compressive relationship between the physical stimulus (amplitude) and average firing rate of auditory cells. And or psychophysical judgments of sound intensity directly correspond to these firing rates.

So much for our brief look at the science of sensation. While many philosophers focus narrowly on one sense-modality, we have considered several. So we can see the “big picture” that emerges. 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 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. The science is *apparently* at odds with tracking intentionalism. But can this be definitely shown? In the following sections, I will construct arguments that are intended to do exactly that.

3. First Argument: The Internal-Dependence Argument

I call my first argument the *internal-dependence argument*. The aim is to *demonstrate* the conflict between tracking intentionalism and science by describing *counterexamples* to tracking intentionalism.

3.1 Why Actual Cases Fall Short

Now you might think that *actual cases* involving perceptual variation suffice as counterexamples to tracking intentionalism. Many philosophers have certainly thought so, including Ned Block (1999), Brian McLaughlin (2003), Sydney Shoemaker (2000), and Uriah Kriegel (2009). But their arguments have been ineffective. Before turning to

my own counterexamples, it will be helpful to see why. For my counterexamples will be designed to preempt the usual responses.

You are probably familiar with actual cases of perceptual variation. The same bodily disturbance-type can give different individuals slightly different pains. The same substance can taste differently to different individuals. There are, for instance, “supertasters”. The same odor cloud can smell differently to different people. The same sound can sound differently. There are individual differences in pain intensity in response to the same noxious stimuli. As Block (2010) has noted, in some cases phenomenal variation can even be due to differences in *attention*. All of these cases pose the same problem for tracking intentionalism. In these cases, the individuals involved have *different* experiences of the same stimulus, presumably due to differences in the kind of neural processes discussed in the previous section. But don’t we have to say that experiences accurately represent the *same* external physical properties? In that case, phenomenal character is determined by more than the external physical property represented in the world, contrary to tracking intentionalism.

In response to an actual case of variation, tracking intentionalists can always invoke what I shall call the “illusion response” (Tye, Byrne and Hilbert, Batty) or else the “pluralist response” (Kalderon, Smith).⁵

The *illusion response* is especially plausible in cases where there is some kind of interference or abnormality, so that *optimal conditions* do not obtain. The idea is that the individuals’ different experiences *do not* represent the same external physical property of the stimulus. One individual represents a physical property that the stimulus does have and the other represents a physical property that it does not have. This representational difference constitutes the phenomenal difference.

The tracking intentionalist might invoke *pluralist response* in cases where optimal conditions obtain, and the individuals involved are normal. The idea is that the individuals are actually *tracking different physical properties* of the stimulus, because of differences between their sensory systems. So the sensible qualities they represent are distinct but equally real properties of the stimulus. They both get it right. This response is “pluralist” because it says that the external world is rich with sensible qualities. This helps the externalist explain perceptual variation without having to posit illusion. So, for instance, a wine might

⁵ For the illusion response to some cases, see Byrne and Hilbert (2003), Tye (2006), Batty (2009). For the pluralist response to some cases, see Kalderon (2011) and Smith (2007, p. 65).

actually have many objective tastes, constituted by overlapping but distinct chemical types. And, due to differences in their taste systems, one individual might perceive one while another individual perceives another (Smith 2007, 65). Indeed, to handle more radical cases of *interspecies variation*, the tracking intentionalist will say that foods have various alien tastes that we cannot imagine. The human tastes and the alien tastes are constituted by different chemical properties belonging to the same substances. We track and thereby perceive one range of properties of the substances. Another species might track and thereby perceive a different a totally different range of properties of the substances.

The idea here is that actual differences in neural processing between individuals lead to differences in what external properties they track and thereby represent. In this way, the tracking intentionalist can handle actual cases of variation.⁶

To refute tracking intentionalism, what we need is a case that is invulnerable to both the illusion response and the pluralist response. At this point, some philosophers might be tempted to invoke intuitions about the possibility of far out hypothetical cases, like brains in vats and spectrum inversion. For instance, since there is an explanatory gap between color experience and the reflectance (light-involving) properties of surfaces, it is quite conceivable that two individuals should accurately track the very same reflectance property but have different color experiences (Shoemaker 1994). But tracking intentionalists just reply that, while this may be conceivable, it is not possible (Tye 2000). Indeed, I would add that such intuitions are just instances of our more general explanatory gap intuition to the effect that experience is modally independent of *all* physical conditions. So physicalists have special reason to be suspicious of them.

⁶ But I think that some extreme cases of variation, not discussed in the literature, are particularly troublesome for tracking intentionalists and objectivists. As Batty (2009) notes, a large percentage of humans cannot smell *androstenone*. She does not note that, of those 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 smell. Here the illusion response would be implausible, given the parity between the groups. And the pluralist response (Kalderon, Smith) is problematic as well. In one version, the pluralist view would have it that the floral smell perceived by the first group is identical with the disjunction of all the molecular types (including androstenone) that are the objective correlate of the perception of that smell among humans; and the ruinous smell perceived by the first group is identical with overlapping but distinct disjunction of all the molecular types (including androstenone) that are the objective correlate of the perception of that smell among humans. On this view, androstenone objectively possesses two *radically different* smells. This is hard to accept.

My strategy will be quite different. I will describe hypothetical but realistic *coincidental variation cases*. In these cases, there is neural and behavioral variation between the members of different species. Nevertheless, I will simply *stipulate* that, whatever conditions need to be in place in order for two creatures to accurately represent exactly the same properties, those conditions are indeed in place. While there is neural and behavior variation, there is a *complete coincidence* in what objective properties their sensory systems track. Given the vast neural and behavioral differences, I will argue that they would have different experiences. The cases are not just ones of alternative “neural realizations” of the same experience. To establish this verdict I will *not* use dubious intuitions which tracking intentionalists might simply dismiss; I will *argue* for this verdict on the basis of the research in neuroscience and psychophysics discussed previously. But tracking intentionalism (and indeed all versions of externalist intentionalism) delivers the mistaken verdict that the individuals have exactly the same experiences. Given my stipulations, neither the illusion response nor the pluralist response will be available to externalists.

I will consider several cases. By focusing on several cases, we can appreciate the strength of the cumulative case against the externalist program. To answer my argument, externalists would need to develop solutions in every case, instead of narrowly focusing (as they often do) on one sense-modality.

3.2 *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. It is reasonable to think that while the berries taste horribly bitter to us, they taste different to the birds.

Now this actual case is no problem for tracking intentionalists. They can appeal to the pluralist response. Humans and the birds differ at the receptor 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 between the humans and the birds is grounded in their sensorily representing different, but equally real, taste properties of the berries. In short, they can invoke the *pluralist response*.

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 is still reasonable to think that berries taste differently to them, but tracking intentionalism is inconsistent with this verdict.

In more detail, suppose *Yuck* and *Yum* belong to different species that evolved in separate environments containing some berries. 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. By contrast, 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 activation states (spatiotemporal neural patterns discussed in §2). Yuck and Yum 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.

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 state 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 (see Figure 4). And there are consequent differences in their behavioral responses.



Figure 4. Yuck and Yum's neural representations of the berries occupy different locations in the kind of neural similarity space uncovered by Smith and coworkers (1983). Adapted from Churchland (1996) with permission.

Despite these differences, we can stipulate that Yuck and Yum are similar at the receptor level. Indeed, we can stipulate that, when they taste the berries, the *postreceptor* ensemble activation patterns in their taste systems, although different, optimally track the very same complex chemical property of the berries, *C*. This chemical property *C* will likely be a *disjunctive property*, because many different combinations of chemical properties can produce the same response in the taste system. So I am stipulating that their ensemble activation states track the same *disjunction* of chemical properties *C*, the very one with which tracking intentionalists and other objectivists about taste would identify the taste perceived by Yuck and Yum.

That, then, is the case described in non-phenomenal terms. The crucial question 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 are 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. 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 bitter-tasting things, whereas Yum's taste experience of the berries is probably similar to his taste experience of bananas and other characteristically sweet-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 are not behaviorists or functionalists, this suggests that they have different experiences of the berries, just as humans and birds do in the actual world.

Note that I only say that, given what we know about the physical basis of taste experience, it is reasonable to suppose that there are *some* 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*. The functional, sensorimotor differences may play a role too.

So the only reasonable verdict is that Yuck and Yum would have different taste and smell 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.

By stipulation, Yuck and Yum's ensemble activation states, although different, track exactly the same complex 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, this means that Yuck and

Yuck and Yum both sensorily represent (“perceive”) the very same chemical property, *C*, of the berries. In general, their ensemble activation states, although different, represent exactly the same chemical properties.

Consider an analogy. The words ‘river’ and ‘liver’ are similar but represent quite different things. Likewise, on tracking intentionalism, even though Yuck and Yum’s ensemble activation states occupy different locations in neural similarity space, they represent the very same external chemical properties: on this view the similarity/difference structure of those states is incidental to what they represent.

So far I have just argued that tracking intentionalism implies that Yuck and Yum’s ensemble activation states are representationally identical. Now I do not think this is *by itself* bad. (Here I am responding to a query from Fred Dretske.) I agree with tracking intentionalists like Dretske and Tye that 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. In particular, tracking intentionalists claim that the *sensory character* of taste experience (sweet, bitter, etc.) is determined by what chemical properties taste experience represents. Once we add this to the tracking intentionalist’s commitment to the claim that Yuck and Yum’s experiences represent the very same chemical type, we *do* get a bad verdict: that Yuck and Yum have experiences of the berries identical in sensory character (sweet, bitter, etc.), despite the vast neural and behavioral differences between them. For instance, on tracking intentionalism, maybe they both have an *intensely bitter* experience of the berries. *This* verdict is radically implausible. Their ensemble activation states occupy different locations in neural similarity space; and that is the only known predictor of taste quality (sweet, bitter, etc.). And, by stipulation, Yum is drawn to the berries. So by far the most reasonable verdict is that Yuck and Yum would have experiences of the berries (Doll’s-eyes) that differ radically in their sensory character (bitter and sweet), just as humans and birds do in the actual world.

Barry Smith (2007) is an objectivist about tastes who handles actual cases of taste variation by accepting the pluralist view that foods have

multiple tastes that different individuals can perceive depending on the conditions. But he cannot handle Yuck and Yum similarly. On his view, the taste that Yuck perceives is an enormously complicated property of the berries. (Smith would say that *flavor* is an even more complex property, because unlike taste flavor depends on odor; perhaps flavor is a configuration of sapid, odorous and textural properties tracked through several sense organs. Here I am focusing on taste, but my argument applies *mutatis mutandis* against objectivism about flavor.) Yuck's perceiving the taste presumably supervenes on his bearing some natural relation (perhaps a tracking relation) to it. But I stipulate that Yum bears the same relation to the very same enormously complicated property of the berries. So, even if Smith is right that the berry has multiple objective tastes, he is committed to the claim that Yum perceives very same one as Yuck. The trouble is that this is implausible. It is not the case that they perceive the same taste (as it might be, an extremely bitter one), which Yuck does not like and Yum happens to like. Given that they differ in the kind of ensemble activation states that we know to the best predictors of sensory quality and sensory similarity, and given that they differ in their fine-grained taste-related behavior, the only reasonable verdict is that they ostensibly perceive quite *distinct* tastes (as it might be, a bitter one and a sweet one). The conclusion I draw is that tastes are not complex objective properties of external items.

3.3 Sniff and Snort

In the previous section, I discussed an fMRI study by Howard *et al.* on humans. Let *Sniff* be one of the actual participants of the study. He successively smells the chemical structures represented in Figure 1. Among other things, he reports that his (+R) *limonene* experience resembles his *citral* experience more than his *menthol* experience. In particular, limonene as well as citral presents a citrus smell, whereas menthol presents a mint smell. Why is this?

The explanation is not to be found in the (possibly disjunctive) chemical properties – call them *L*, *C* and *M* – that his smell experiences track, and with which the objectivist would identify the smell qualities perceived by Sniff. There is no evident sense in which limonene is more like citral than menthol. In fact, if anything, limonene is more like *menthol* than citral. This is an actual case of bad external correlation.

Howard *et al.* claim that the explanation is only to be found in Sniff's brain, in particular, in his posterior piriform cortex (PPC). The explanation for why Sniff's limonene experience resembles his citral experience more than his menthol experience is that Sniff's PPC *neural representation* of limonene is more similar to his *neural representation* of citrus than in his *neural representation* of menthol. As Margot (2009, 814) says, "the PPC codes [i. e. determines] odor quality rather than structural and chemical similarity [in the odorants tracked]".

This appears directly at odds with the externalist account of sensory quality promoted by tracking intentionalists. To make the conflict precise, consider a counterfactual situation. In this situation, everything is more or less the same as in the actual world. But, due to some chance differences in our evolutionary history, there are differences in the postreceptor wiring leading from the receptors in the nose to ensemble activations further downstream in the PPC. In this situation, Sniff's counterpart, *Snort*, participates in the Howard study. Because of the difference in wiring, while Sniff's neural representation of limonene is more similar to his neural representation of citrus than in his neural representation of menthol, Snort's neural representation of limonene is more similar to his neural representation of *menthol* than his neural representation of citral. In short, their PPC representations of limonene occupy quite different positions in the kind of *neural similarity space* studied by Howard *et al.*

Let us suppose that there are consequent behavior differences. So, while Sniff sorts limonene with citral and not menthol, Snort sorts limonene with menthol and not citral. Since Snort evolved in a different counterfactual situation, he probably does not speak a language that looks like English. But we can suppose that, while Sniff reports that his limonene experience resembles his *citral* experience more than his menthol experience, according to the best translation Snort says "my limonene experience resembles my *menthol* experience much more than my citral experience".

Of course, the most reasonable verdict about this case is that there are phenomenal differences between Sniff and Snort's smell experiences of limonene, citral and menthol. In particular, their smell experiences fall into different phenomenal resemblance-orders. And while limonene presents a citrus smell to Sniff, it presents a minty smell to Snort. For Sniff and Snort's PPC *neural representations* fall into different resemblance orders. And *neural* similarity is the only thing in the physical world known to predict *smell* similarity. In addition, this verdict is the

best way of making sense of Sniff and Snort's different sorting and verbal behaviors.

But, if we add some more details to the case, tracking intentionalism delivers the opposite verdict that there are no phenomenal differences among Sniff and Snort's smell experiences of limonene, citral and menthol, despite the radical neural and behavioral differences between them. For we can stipulate that their neural representations for these chemicals, while different, track under optimal conditions the very same chemical properties of those chemicals, namely them *L*, *C* and *M*. On tracking intentionalism, the smell qualities they perceive are none other than these chemical properties. So, tracking intentionalism, Sniff and Snort represent and hence perceive exactly the same smell qualities. In general, on tracking intentionalism, the representational contents of their experiences are *exactly* the same across the actual case and the counterfactual case, if the tracking facts are held constant. There are differences in the "neural content carriers" but no differences whatever in the externally-determined representational contents they carry. Because there are no representational differences between Sniff and Snort, there also can be no phenomenal differences between them, according to intentionalism. *A fortiori*, there can be no differences in the phenomenal *resemblance-orders* of their experiences of limonene, citral and menthol. And, on tracking intentionalism, Snort perceives the same specific objective "citrus" quality in limonene that Sniff perceives. This is so despite the radical neural and behavioral differences between them

In reply, the tracking theorist might grant that on his theory Sniff and Snort track and thereby represent the odorants as having the same *monadic* olfactory properties, namely *L*, *C* and *M*. But he might insist also that they represent these same properties as standing in *different resemblance-orders*. In particular, Sniff's odor experience represents *L* as more like *C* than *M*, while Snort's odor experience represents *L* as more like *M* than *C*. In this way, the tracking intentionalist can insist that there is a representational difference, and hence a phenomenal difference, between their experiences. Call this the *structure gambit*, because the idea is that their experiences represent different contents about qualitative structure.

But this response fails for a number of reasons. Here I will only mention three. First, in the Howard experiment Sniff and Snort smell the odorants *successively*. So, contrary to the reply, there is no time at which their olfactory experience might have represented them (or the

general properties *L*, *C* and *M*) as standing in a certain resemblance-order. Analogy: if you experience three objects *successively*, your experience cannot represent a spatial relation among them. I take this point from Alex Byrne (2003, 656). Second, in any case the present reply is inconsistent with tracking intentionalism. I stipulate that Sniff and Snort bear the tracking relation (and other relevant relations) to the *same* conditions. So even if we concede for the sake of argument that Sniff tracks and thereby represents a certain complex structural condition, then Snort represents the *same* structural condition. There is no naturalistic account of how they might represent *different* such conditions. Third, on the “structure gambit”, Sniff and Snort perceive exactly the same individual smell qualities, but experience them as standing in different resemblance-orders. But this can be ruled out *a priori*. It is like saying that, while you actually experience blue as more like purple than green, another creature could experience blue as more like *green* than purple. Against this, perceiving individual qualities essentially involves perceiving them to stand in certain resemblance relations.

Clare Batty (2009) is partial to objectivism about smell qualities, even if she does not outright endorse it. She suggests that it can survive arguments based on actual cases of variation. I believe that the hypothetical “coincidental variation” case involving Sniff and Snort poses a difficulty for any version of objectivism that is more serious than the problem she discusses concerning actual cases. Here is the argument. According to the objectivist, when Sniff smells limonene, the smell quality that he perceives is identical with some “objective” property of the odor cloud. Call it *L*. Maybe *L* is a disjunctive chemical property, where the disjuncts are all the combinations of chemical properties that yield that same distinctive smell. The details do not matter. Now, although she does not discuss this issue, Batty (if she outright accepts objectivism) must say that Sniff perceives *L* by virtue of bearing some (perhaps as yet unknown) complex physical relation *R* to *L*. Presumably, *R* would a kind of tracking relation, but my argument is neutral here. Let us stipulate that in the counterfactual situation, on smelling limonene, Snort also bears *R* to the same complex response-independent property *L*, despite his neural and behavioral differences from Sniff. In general, Sniff and Snort both bear to this same property *L* all the naturalistic relations that might ground perceptual representation, despite the neural and behavioral differences between them. Then the objectivist is committed to the claim that on smelling limonene Sniff and Snort perceive the very same smell quality, despite the PPC

neural differences and behavioral differences between them. But this is very implausible. The PPC neural differences and behavioral differences between Sniff and Snort make it reasonable to suppose that they ostensibly perceive *distinct* smell qualities. (They don't perceive the same smell quality under different modes of presentation, whatever that might mean; they perceive *numerically distinct* smell qualities.) In particular, Sniff ostensibly perceives a citrus smell quality and Snort perceives a minty one. The conclusion I draw is that smell qualities cannot be objective properties of odor clouds.

3.4 *Mild and Severe*

As I mentioned in §2, the relationship between stimulus intensity and pain intensity is generally one of “response expansion”. Pain intensity also depends on stimulus size and duration, in a complex way. In short, there is bad external correlation. For instance, in one experiment, Donald Price (1999) asked subjects to rate their sensory pain intensity on a visual analogue scale (VAS) with a sliding marker in response to noxious temperature. He consistently found that the psychophysical relationship of pain sensation intensity to heat stimuli (45–50°C, for five seconds) is a positively accelerating power function with an exponents of about 3.0. Of course there might be indeterminacy concerning the precise number. In any case, small changes in temperature yield large changes in perceived pain intensity. This makes evolutionary sense: small changes in stimulus intensity can be very dangerous. For instance, in another experiment, Price found that when subjects perceive a standard stimulus of 47°C, they determined 49.8°C (average) to be roughly “twice as intense” as 47°C, which agrees perfectly with the prediction of his independently-determined stimulus-response curve with an exponent of about 3.0 (see Figure 2 in §2).

Two clarifications. First, we must bear in mind that pain researchers distinguish between the sensory dimension of pain and the affective (unpleasantness) dimension of pain. These can come apart. The same pain can “bother” different people to different degrees. The results I have described concern subjects' ratings of the *sensory* intensity of pain. Price also asked subjects to rate the *unpleasantness* of their pains, and reliably found different results. But this will not concern us at present since I will be concerned with *sensory* intensity (for more on this see §5). Second, you might think the notion of one pain being twice more intense than another makes no sense. But ratio scaling of loudness makes sense, so why not pain intensity? Of course there might be

a great deal of indeterminacy. But why can't it true that one pain is *roughly* twice greater than another in intensity? Price provides evidence that rough claims like this are true. However, the argument I am about to present against tracking intentionalism does not strictly speaking require the truth of these particular judgments; even if the judgments are not true, the fact that subjects make them can provide evidence about what their experiences are like.

While pain intensity is related in a non-linear fashion to numerous stimulus features (bad external correlation), it is more proportional to firing rates in the brain (good internal correlation). For instance, as Price (2002, 395) reports, numerous studies have found that "stimulus-response functions of WDR [wide dynamic range] neurons to this range of skin temperatures are precise positively accelerating power functions, ones that strongly resemble power functions that are obtained from human subjects who rated these same temperature stimuli" (see for instance Figures 3 and 2 in §2). Indeed studies have shown that *sensory intensity* of pain is linearly related to those neural firing rates, as we saw in §2.

All of this goes against tracking intentionalism; indeed it goes against any "externalist" theory of pain. Contrary to this view, the intensity of a painful response to temperature is not merely determined by any features of the stimulus that our pain system tracks. The simplest hypothesis is that it is more directly determined by S1 firing rates. To turn this into an argument into tracking intentionalism, all we have to do is describe a coincidental variation case involving pain.

Let us consider a counterfactual situation in which the psychophysical response curve describing the relationship between noxious temperature and neural response in S1 and other neural regions is *steeper* than it is in the actual world – that is, steeper than those shown in Figures 3 and 2 in §2. In other words, the rough exponent is consistently much higher than it is in the actual world (than is, higher than around 3.0).

Let us suppose that this is not because in the counterfactual situation the same noxious temperatures are more of a threat to the organism, or more likely to jeopardize the organism, than they are in the actual world. My idea is that it just happens that in this situation the psychophysical response curve is steeper than in the actual world. While selection pressures might ensure that the response function has an exponent greater than 1, so that we will be sure to avoid increasing noxious temperatures, the precise value (or range of values) of the ex-

ponent is evidently a matter of chance. So it can vary across worlds in which noxious temperatures are equally dangerous. (I mention this because it will later on be important to undermining proposals suggested by Cutter and Tye and Hill.)

Now consider a subject of one of Price's experiments in the actual world, where the response function is *less steep* that it is in the counterfactual situation. Call him *Mild*. He experiences temperatures in the noxious range between 45°C and 50°C and undergoes various S1 firing rates. In consequence, he rates his pains using the VAS scale. In addition, he judges 49.8°C to be roughly twice as intense as 47°C.

Now consider Mild's counterpart in the counterfactual situation. Call him *Severe*. Like Mild, Severe experiences noxious temperatures between 45°C and 50°C and undergoes increasing S1 firing rates. However, because in this situation humans' psychophysical response function is steeper, Severe's S1 firing rates in response to these same temperatures are much higher than Mild's. There is not just a difference in absolute firing rates; Severe's S1 firing rates increase *more rapidly* than Mild's. So, for instance, while moving from 47°C to 49.8°C roughly doubles Mild's average S1 firing rate, it far more than doubles Severe's S1 firing rate. In consequence of these neural differences, there are also behavioral differences between Mild and Severe. Using the VAS scale, Severe consistently rates his thermal pains as more intense than does Mild. While Mild reports 49.8°C to be twice as intense as 47°C, Severe reports 49.8°C to be "much more than twice as intense" as 47°C. (Since Mild has a different evolutionary history than Severe, it is unlikely he speaks a language that sounds like English; but suppose that this is the best translation of his report.) Finally, Severe responds to 49.8°C with much greater urgency than does Mild; his pulse rate is higher; and so on.

Despite these neural and behavioral differences between Mild and Severe, let us stipulate that their pain states track exactly the same properties of the thermal stimuli. In general, whatever kind of relations the tracking intentionalist thinks ground representation (whether they be Tye's simple tracking relations, or more Millikan's more complex teleological relations), Mild and Severe bear those relations to exactly the same properties of peripheral stimuli.

So far I have described Severe's situation in neural and behavioral terms. I haven't stipulated anything about his pain experiences. That is the crucial issue.

In my view, Severe would have pains that are more intense than Mild, in response to the same noxious thermal stimuli. There are differences between them at the sensory level, not just the affective level. So for instance, the difference in intensity between Severe's pains at 49.8°C and 47°C is greater than the difference between Mild's pains at those intensities. The case for this is obvious. First, the relationship between noxious temperature and S1 firing rate is steeper in Severe than in Mild, and S1 firing rate is the best-known predictor of the sensory dimension of pain. (Recall that S1 activity codes for the sensory dimension of pain, while the ACC codes for the affective dimension.) Second, the psychophysical and other behavioral differences between Severe and Mild are only explained by the claim that Severe's pains are more intense than Mild's on the sensory dimension.

But, of course, tracking intentionalism delivers the opposite verdict. Indeed, the problem here is not just a problem for self-described "tracking intentionalists" like Tye and Dretske. It is a problem for "externalist intentionalists" in general. So, for instance, Hill does not advocate a specific theory of representation, although he does sometimes appeal to a kind of tracking theory (2009, 149, n. 16; 179–80). Like Tye and Dretske, he is somewhat undecided on the issue of exactly what peripheral physical properties various types of pain represent. But, whatever kind of relations the externalist thinks ground representation (whether they be Tye's simple tracking relations, or more Millikan's more complex teleological relations), it seems that we can stipulate that Mild and Severe bear those relations to exactly the *same* properties of peripheral stimuli. In that case, the externalist is committed to saying that they have phenomenally different experiences. Here he cannot appeal to the illusion response or the pluralist response (§3.1). I have simply *stipulated* that, whatever conditions need to be in place in order for two creatures to accurately represent exactly the same properties, those conditions are indeed in place in the case of Mild and Severe.

Cutter and Tye (2011) and Hill (2012) have offered a response to cases like this. Applied to the present case, the idea would be that 49.8°C is more of a "threat" to Severe than it is to Mild; in particular, it has the property *being bad to degree D to Mild* and the property *being bad to degree D* to Severe*. And Severe *somehow* represents the first "valuational" property while Severe represents the second. Because their experiences represent different "valuational properties", the tracking intentionalist can say that they differ at least on the affective dimension

of pain, even if he must say that they have exactly the same sensory intensity. But this response fails to apply to my present case. (i) As I have set up the case, 49.8°C is simply *not* more of a threat to (“more likely to harm”) Severe than it is to Mild, as I explained above, so here the valuational response cannot get off of the ground. (ii) In any case, the valuational response does not fit with the science. While S1 activity codes for sensory intensity, ACC activity codes for affect. And Mild and Severe differ in *both* S1 activity and ACC activity. Moreover, their VAS ratings for *sensory* intensity differ. So the only reasonable verdict is that their pains differ on the sensory dimension, not merely the affective dimension. I will have more to say about these issues in §5.

3.5 *Soft and Loud*

My final counterexample to tracking intentionalism concerns the perception of loudness. Tracking intentionalists are objectivists about loudness and other audible qualities, claiming that they objective properties of sound events. My final hypothetical case poses a problem for *any* view that incorporates objectivism about audible qualities, not just tracking intentionalism. Since Casey O’Callaghan (2002) has provided the most sophisticated defense of objectivism about auditory qualities, I will focus on his approach.

To begin with, O’Callaghan holds that particular sounds, the *bearers* of audible qualities, are events of oscillating or interacting bodies disturbing or setting a surrounding medium into wave motion. He holds that auditory qualities are objective, physical properties of these physical events. In the case of loudness, which will be my focus, his basic view is that particular loudness levels are identical with complex properties of these sound events, properties involving amplitude, frequency, “critical bands”, and duration. A complex account is needed, because while loudness is mainly related to amplitude it also depends on other parameters. So, for instance, the loudness of a 40 dB pure tone at 500 Hz is the same as that of a 60 dB pure tone at 50 Hz. This is represented in so-called equal loudness curves. It is analogous to metamerism in the domain of color vision: many different combinations of lights can yield the same color perception. So, if we confine ourselves to pure tones, a given loudness level might be identified with the *disjunction* of all the different amplitude-frequency pairs that give rise to a perception of that loudness level. The account would have to be even more complex than this, because loudness also depends on critical bands and duration. In any case, the point is that the objectivist will not identify

loudness levels with simple amplitudes or intensities; he will identify them with much more disjunctive, unnatural properties. As O'Callaghan says, these properties will only be of "anthropocentric interest", because they are the objective correlates of human loudness perception.

Now suppose that Soft is an actual human who hears tones of increasing amplitude. As mentioned in §2, there is a non-linear, highly compressive relationship between amplitude and perceived loudness level (holding frequency fixed). Therefore, huge differences in amplitude are needed to generate small differences in perceived loudness. For instance, he judges tone B to be twice louder than tone A when the amplitude of B is ten times louder in intensity (which is related to amplitude). The explanation for why there is a compressive relationship between amplitude and perceived loudness is that there is a matching compressive relationship between amplitude and total neural activity in the auditory channel, as several experiments have shown. As we saw in §2, there is evidence that perceived loudness is more proportional to total neural response (even in the cortex) than it is to amplitude and other complex physical parameters. While there is bad "external correlation", there is evidence for better "internal correlation".

Now consider a counterfactual situation in which humans evolved so that the same auditory stimuli normally produce a *greater* total neural response in the auditory channel than they do in the actual situation. In this situation, Soft has a counterpart, Loud. When Loud hears the same tones that Soft hears in the actual situation, the total neural response in his auditory channel increases *more rapidly* than Soft's. (So the case is similar to that of Mild and Severe.) In consequence of these neural differences, there are also behavioral differences. Thus, in this counterfactual situation, humans' subjective estimations of loudness yield stimulus-response functions that are much steeper than those which characterize the loudness perception of actual humans. For instance, Loud and other normal perceivers consistently report that tone B is *much more* than twice as loud as tone A. In addition, Loud and other normal perceivers in this counterfactual situation are more likely than actual humans to *notice* the same auditory events; amplitudes that do not produce discomfort in actual humans, produce discomfort in humans in the counterfactual situation; and so on. Why did humans in this situation evolve auditory systems that differ from our own in amplifying the neural responses responsible for loudness perception? It does not really matter to my argument. Maybe they rely on hearing

more than we do; or maybe they evolved in an environment in which they must notice certain sounds more readily.

So Soft and Loud differ in their loudness-related neural and behavioral responses. Yet I also want to stipulate that there is a complete coincidence in the objective auditory properties they track. As Soft hears the tones in the actual world, he perceives increasing loudness levels. As we saw, O'Callaghan identifies these perceived loudness levels with complex, probably disjunctive physical properties of the tones: $D1, D2, D3, \dots$. Now, although O'Callaghan does not address this issue, he must hold that Soft perceives, or sensorily represents, $D1, D2, D3, \dots$, because his cortical neural representations (those which realize his auditory experiences) bears *some* naturalistic relation R to $D1, D2, D3, \dots$. Maybe it is a kind of tracking relation; or maybe it is difficult to specify, because providing a theory of perception, or perceptual representation, is difficult. But there must be some naturalistic facts that determine that Soft perceives $D1, D2, D3, \dots$ to the exclusion of all of the other candidates. Now, whatever the naturalistic relation R is, I stipulate that Loud and Soft's corresponding cortical neural representations bear to R to the same properties, namely $D1, D2, D3, \dots$. True, Loud's neural representations involve higher neural firing rates than Soft's, and result in different behavioral responses. But the stipulation here is that they nevertheless track or detect the same objective properties $D1, D2, D3, \dots$ of the sound-events. Compare: in different types of mercury thermometers, different mercury heights can track the very same objective temperatures. Therefore, whatever the relation R is, the stipulation I am making is apparently possible, and we would need a good reason to think otherwise.

To clarify, my stipulation here is not that Loud and Soft's neural representations bear relation R to the same simple *amplitudes* or *physical intensities*. According to O'Callaghan, loudness levels are not mere amplitudes or intensities. Instead, he maintains that they are the more disjunctive, complex properties $D1, D2, D3, \dots$ which involve not just intensities but also frequencies and "critical bands". What I am stipulating is that Loud's neural representations, as well as Soft's, bear relation R to *these* properties, the very properties with which O'Callaghan identifies the (low) loudness-levels perceived by *Soft* in the actual world.

Now you can see how this creates a problem for O'Callaghan's objectivism about audible qualities as well as tracking intentionalism about the perception of audible qualities. Given the vast neural and

behavioral differences between them, together with what we know about the physical basis of loudness perception, the most reasonable view is that Soft and Loud auditory experiences differ phenomenally as regards intensity. Given this, it follows that they ostensibly perceive *different* loudness levels, when presented with the same tones. In particular, when they hear the same sequence of tones, Loud perceives *higher* loudness levels than Soft, which increase more rapidly than those perceived by Soft. It is not the case that they perceive the same loudness levels (properties) via different “modes of presentation” or “mental paint”, whatever that means; the correct description of the phenomenal difference is that they ostensibly perceive distinct loudness levels (*pace* Block 2010, p. 25). But, given O’Callaghan’s objectivist theory of loudness, together with the natural assumption that perception must be grounded in some naturalistic relation *R*, it follows that Loud perceives *the very same* loudness levels as Soft, which he identifies with *D1, D2, D3, . . .* The case also undermines tracking intentionalism about auditory experience. For, on tracking intentionalism, they have phenomenally identical auditory experiences, despite the vast neural and behavioral differences between them, because their experiences “represent” exactly the same objective properties. I conclude that tracking intentionalism is mistaken. Indeed, we must reject any view on which audible qualities are *objective* properties like *D1, D2, D3, . . .* Maybe they are response-dependent properties. Or maybe they are internal neural properties “projected” onto external sound-events. But they are not *objective* properties like *D1, D2, D3, . . .*

Now you might wonder why *actual* cases of auditory variation are not enough to refute tracking intentionalism and indeed any objectivist theory of audible qualities. The reason is that proponents of such views can appeal to the pluralist response or the misrepresentation to handle actual cases (§3.1). By contrast, in my hypothetical case of Soft and Loud, I have simply *stipulated* that, whatever conditions need to be in place in order for two creatures to accurately represent exactly the objective audible qualities, those conditions are indeed in place. So, in this case, neither the pluralist response nor the illusion response is available.

O’Callaghan (2009, sect. 3.2.5) considers a partial error theory as a response to “bad external correlation”. The idea is that, in the actual world, auditory experiences (accurately) represent individual audible qualities, but also “distort their magnitudes of difference”. The proponent of this view might grant that his view entails that Soft and Loud’s

experiences represent the same individual loudness levels $D1, D2, D3, \dots$ on hearing the tones. But he might insist that their experiences (inaccurately) represent those same loudness-levels as standing in different *magnitude relations*, which accounts for the phenomenal difference. This is a version of what I called the “structure gambit” in connection with the case of Sniff and Snort. For a few reasons, it is untenable. (i) There is Alex Byrne’s point (2003, 656). Since Soft and Loud hear the tones and their apparent loudness-levels successively, and not at the same time, there is simply no time at which their *experiences* might represent all of those loudness-levels as standing in different magnitude relations. (ii) In any case, since Soft and Loud are exactly alike in their relevant naturalistic relations to the environment, there is no obvious naturalistic account of how Soft and Loud’s experiences might represent *different* contents involving qualitative structure. (iii) This sort of view can be ruled out *a priori*. For instance, if Soft hears three loudness levels as increasing by equal intervals, then Loud cannot hear the *same* loudness levels as increasing by greater magnitudes. The right description of the case is that Loud is hearing *different, higher* loudness-levels than Soft (which is something that objectivists about loudness cannot accept, as we have seen).

4.4 *The official internal-dependence argument*

Of course, such cases could be multiplied indefinitely. To refute tracking intentionalism, and the general objectivist treatment of the sensible qualities, 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.

Let me make two clarifying remarks. First, recently Michael Tye (2009, 194) has claimed that no empirical work on the explanatory underpinnings of phenomenology can establish the strong internalist claim that microphysical duplicates metaphysically necessitates total phenomenal duplication. But my argument does not depend on the strong internalist claim that microphysical sameness necessarily guarantees phenom-

enal sameness. It only depends on internal-dependence: internal factors play a role, in the very minimal sense that, in *some* coincidental variation cases, the right verdict is Different Experiences. This weak claim is supported by empirical work; and it is enough to refute tracking intentionalism, and indeed (as we shall see in §5) any version of “externalist intentionalism”.

Second, the internal-dependence argument also does not depend on any theory of sensory character. For instance, it does not depend on the claim that experience-types are necessarily identical with neural-types in the head, although it might naturally suggest that view. It also does not depend on the somewhat strange view that tastes, smells, sounds and pains are literally in the head. My case for internal-dependence only relies on the empirical findings and is neutral on the philosophical interpretation of those findings. For instance, since the individuals in coincidental variation cases differ in functional and sensorimotor respects, it is also compatible with functionalist and sensorimotor approaches.

Of course, there are potential objections to the internal-dependence argument; but before addressing objections, I would like to put my second empirical argument on the table.

4 Second Argument: The Structure Argument

In developing my internal-dependence argument, I used good internal correlation as well as bad external correlation to support internal-dependence, which is a claim about non-actual coincidental variation cases. My structure argument is a totally independent argument. It only depends on bad external correlation, which is well confirmed in psychophysics. And it concerns actual cases. Recall that some of the individuals in my coincidental variation cases were actual individuals. The structure argument is meant to show that, given bad external correlation, their judgments concerning *phenomenal structure* in the actual world come out false. It poses a problem for all objectivists about the sensible qualities. Since tracking intentionalists are committed to objectivism, it undercuts their view.

As with the internal-dependence argument, I will illustrate the structure argument by focusing on a few cases. Given the broad similarities across many sense-modalities, if objectivism is true about one type of sensible quality, then it is true of other types. So the objectivists are committed to a general view, even if they often focus on a single case. By considering a few cases, we will be able to appreciate the cu-

mulative case against their view. I will start with an initial, *prima facie* challenge; the real argument will come afterwards, when we look at potential responses.

4.1 Three Illustrations of the Initial Challenge

In the fMRI experiment conducted by Howard (§2), actual subjects made introspective reports on their smell experiences. Suppose that Sniff is one of them and that he makes a report along the following lines:

Sniff's report: The limonene smell quality I experienced is overall more like the citral smell quality than the menthol smell quality.

The structure argument is simple. This introspective report is true. It was just obvious to Sniff. Indeed, to all normal humans, limonene and citral in fact have slightly very similar citrus smells while menthol has a quite different minty smell. But, given bad external correlation, tracking intentionalism would appear to entail that the report is *false*.

Here is why the tracking intentionalist seems forced into accepting an error theory. Sniff makes his report under optimal conditions. So on tracking intentionalism Sniff's smell experiences are veridical: the smell qualities he experiences are identical with the actual objective *chemical characters* of limonene, citral and menthol. (Since many chemical structures might yield the same smell qualities, these might be disjunctive chemical characters.) So Sniff's report is true just in case they satisfy the semantic value of the relational predicate 'x is overall more like y than z' as it occurs in his report. But, because this is an actual case "bad external correlation", they apparently do not. To see this, just look at the representation of these chemical types in Figure 1 in §2. It is very hard to see how the chemical characters of limonene, citral and menthol (in that order) could satisfy the semantic value of 'x is overall more like y than z' in the context of Sniff's report. The chemical character of limonene is overall *quite different* from that of citral. Indeed, if anything, the chemical character of limonene is more like that of *menthol* than that of citral. To answer the challenge here, the tracking intentionalists (who think smell qualities are chemical characters) would have to convince us that, despite appearances, 'x is overall more like y than z' has a semantic value in the context of Sniff's report which is indeed satisfied by the chemical characters of limonene, citral and

menthol (in that order). This is what the externalists must do in order to accommodate the truth of that report.⁷

Notice that, because of “good internal correlation”, an alternative internalist or response-dependent theory of smell easily accommodates the truth of Sniff’s reports. On such a view, the smell qualities are not objective chemical characters. On one version, the relevant smell qualities are literally identical PPC neural types in Sniff’s head. On another version, the smell qualities are identical with the dispositions of external odors to produce those neural types. The fMRI study by Howard shows that those neural types *do* stand in the relevant resemblance-order. While there is bad external correlation, there is good internal correlation. So on such theories Sniff’s report is straightforwardly true.⁸

Such cases obviously pose a challenge to all objectivists about smell qualities, not just tracking intentionalists. For instance, Batty (2009) defends the view that particular smells are odor clouds and general

⁷ Many would say that resemblance is always resemblance in respects. They would say that ‘x is overall more like y than z’ has different semantic values in different contexts, because in different contexts different respects of resemblance can be salient and can be weighted differently (Davies forthcoming). My initial challenge to objectivists about smell qualities proceeds in full awareness of these points (see also Pautz 2006, note 4). My challenge to them is to specify the semantic value of ‘x is overall more like y than z’ *in those contexts in which we make reports of smell similarity*, and also show that the molecular types with they identify the smells really do satisfy this semantic value. The fact that there is “bad external correlation” creates a *prima facie* difficulty here. Incidentally, while I would agree that for *particulars* all resemblance is resemblance in respect of various qualities or properties, I would myself reject this claim when *qualities* themselves are concerned. What are the respects in which color hues, or smell qualities, resemble? We draw a blank. This is because qualities themselves (unlike particulars) have no interesting set of (second-order) properties with respect to which they can be similar or different. So, in some cases, when we say quality Q1 is more like Q2 than Q3, we arguably use the predicate ‘x is more like y than z’ to pick out a conceptually primitive comparative resemblance relation, not a relation that we can unpack by citing some context-specific “respects of resemblance”. (Contrary to Byrne (2003), we do not have in mind similarity in “genuine respects”, but a basic kind of similarity that is not similarity *in respects* at all.) If this is right, then it makes it even harder for objectivists about the sensible qualities to answer the structure argument. For in that case, in order to show that our similarity judgments about colors, and smells, and so on, are *true* relative to their strict or face-value interpretations, they would have to show that the corresponding reflectance properties, molecular properties, and so on, satisfy the same conceptually primitive comparative resemblance relation.

⁸ O’Regan (2011, 99) suggests that even such internalist, neural theories of qualitative similarity face a problem. There are different metrics for measuring similarity among neural states. What selects which one is the “right” one? But the proponent of such a theory might claim that our paradigmatic reports about qualitative similarity come out relative to a natural metric for measuring similarity among neural states that can be uncovered by multidimensional scaling (Howard et al. 2009, 396). Of course, a totally different approach would be to provide a *functional* account of qualitative similarity (e.g. in terms of similarities in functional role perhaps, or dispositions to form sophisticated similarity beliefs), but no one has developed a plausible account along these lines.

smell qualities are objective and presumably physical properties instantiated by these odor clouds. Given bad external correlation, how might the objectivist avoid an error theory concerning our resemblance judgments concerning smells?

Sniff's report above concerns resemblances among general qualities or properties as opposed to particular items. But it worth mentioning that we also make judgments about resemblances about particular odors, which we think of as clouds that linger in the air and that we can inhale through our noses. For instance, Sniff will also report that the limonene odor in the air around him is more like the citral odor than the menthol order. This adds to the challenge. On an objectivist account these odors are just collections of limonene, citral and menthol molecules. But, on the face of it, given how different limonene and citral are in their objective chemical characters, there is no obvious sense in which the limonene cloud is more like the citral cloud than the menthol cloud. So on this account it is very hard to see how Sniff's report about the resemblance-order of the odors is true.

There are other impressive cases of bad external correlation involving smell that could be used to illustrate the argument. For instance, –carvone and +carvone are *mirror images*, yet they smell totally different (minty and caraway), because they nevertheless set up totally different ensemble activation states in the brain (§2). So Sniff will report that the smells of –carvone and +carvone are more different (mint and caraway) than the smells of limonene and citral (again, both citrus). But –carvone and +carvone are not *chemically* more different than limonene and citral. Indeed, the opposite is true: limonene and citral are much more chemically more different than –carvone and +carvone. So, on an objectivist theory of smell, it is very hard to see how Sniff's judgments about the resemblances among these odors and their general qualities might be true. Appealing in *some* way to internal factors seems to be the only option (as I will discuss below).

Now let us revisit Soft. Soft is an actual individual in one of the many psychophysical studies on ratio scaling of auditory sensation. He makes the following introspective report:

Soft's report: The apparent loudness of tone B is roughly twice greater than the apparent loudness of tone A.

Now you might be skeptical of such ratio judgments. But such ratio judgments *sometimes* make sense. For instance, we can report on a ratio

relationship among the *apparent* lengths of two lines. It turns out that we also quite good at ratio scaling of perceived loudness. Indeed, there is empirical evidence for the validity of ratio scaling of audible intensity. Those working in psychoacoustics generally think that such introspective reports can be true (Gescheider 1997).

But, again, on tracking intentionalism, Soft's report is apparently false. We can suppose that optimal condition obtain. So, on tracking intentionalism, his auditory experiences are veridical. The apparent loudness levels are identical with actual complex, disjunctive physical properties of the tones - call them *D2* and *D1* - involving intensity, frequency, critical bands, and duration. Therefore, on tracking intentionalism, Soft's introspective report is true just in case *D2* is roughly twice greater than *D1*. In other words, Soft's *experience* of B is twice as intense as his *experience* of A, just in case the *physical properties represented* by those experiences stand in this relationship. But there is no obvious sense in which the disjunctive property *D2* is roughly twice greater than the disjunctive property *D1*, in the way that one length property can be twice greater than another. The main issue here is that there "bad external correlation", in particular, response compression. So, in this case, *D2* involves an intensity that is *much more* than twice greater than that involved in *D1*. As a general rule of thumb, for pure tones, doubling loudness requires increasing intensity as a factor of *ten*. So, tracking intentionalists must *apparently* say that Soft's introspective report is false. Indeed, objectivists about loudness such as Casey O'Callaghan (2002) face the same problem. And the problem arises for other audible qualities. For instance, introspective judgments of equal pitch intervals do not correspond to equal frequency intervals. So on objectivism how can these judgments be true?

Those who advocate internalist or response-dependent theories of auditory quality have less of a problem here. On these views, sound qualities are neural responses or else dispositions to produce neural responses, not objective physical properties of sound-events. As noted earlier (§2), many studies show that neural response (especially in the cortex) is more proportional to auditory intensity than anything in the external world. It's still early days, but maybe it will turn out that Soft's total neural response to tone B is actually twice greater than his total neural response to tone A in terms of overall firing rate. Or maybe it is "twice greater" in some functional sense.

Finally, as we saw, in psychophysical experiments on pain, it has been found that on average merely going from 47°C to 49.8°C doubles

subjects' perceived pain intensity because it vastly increases the neural response in the pain matrix. This is another actual case of bad external correlation; it is a case of response expansion. Thus, in the actual world, a human subject, Mild, will come out with the following report:

Mild's report: My 49.8°C pain is roughly twice more intense than my 47°C pain.

Now if you think ratio scaling of pain makes no sense, and reports like Mild's could not possibly be true, then I could fall back on my previous cases about Soft and Sniff. However in fact that there is evidence that in the right circumstance Mild's report can be true (Price 1999). But, on tracking intentionalism, Mild's report apparently comes out false. On intentionalism, Mild's pains stand in the ratio relation just in case the represented pains do. Since optimal conditions obtain (response expansion is part of the normal function of the pain system), tracking intentionalists must say that the represented pains are the actual peripheral disturbances. Now, in this case, there is no actual tissue damage. So on their view, what are the disturbances in this sort of case? They have not discussed this issue. The physical stimuli, 49.8°C and 47°C? But there is no measure relative to which 49.8°C is twice greater than 47°C. And even if there is one, Mild certainly didn't have it in mind when he made his report. So on this option his report is false. Alternatively, the externalist might say that the represented pains are the peripheral neural patterns, *N2* and *N1*. But it is implausible that Mild's experiences represent peripheral neural patterns, because their function is presumably to indicate potential damage or danger instead (more on this in §5). And even if these are the relevant disturbances, Mild's report might come out false. Maybe *N2* *less than* twice greater than *N1*, and the response expansion (the magnification of the neural response) occurs further downstream.

Again, an internalist theory of pain has less of a problem accommodating our structure judgments. Indeed, in view of good internal correlation and bad external correlation, an internalist theory seems inevitable. Kenshalo's findings (see Figure 3 in §2) indicate that, just as pain intensity doubles between 47°C and 49.8°C in humans, so average WDR neural response in monkeys S1 roughly doubles between these temperatures. And in a fMRI study directly on humans, Coghill et al. found *linear* relationships between pain intensity and neural response to noxious temperatures in S1 and other brain regions. So, while it's

early days, current research suggests that Mild's S1 neural response to 49.8°C might literally be roughly twice greater than his S1 neural response to 47°C. So if pains are S1 neural states or states directly supervenient on S1 activity, Mild's report comes out true.

4.2 Three Unsatisfactory Responses to the Challenge

Of course, what I have said so far only poses an initial challenge to externalists of various stripes. To complete the structure argument, I would have to eliminate all responses to the initial challenge. I will look at the three most obvious responses (some less obvious ones will be addressed in §5).

Error Theory. One response would be to simply accept my argument that tracking intentionalism and objectivist accounts of sensible qualities in general lead to an error theory concerning structure judgments. So, for instance, O'Callaghan (2009, sect. 3.2.5) briefly considers a *partial* error theory about our judgments about magnitude relations among *audible qualities*. In the case of loudness, the idea is that our beliefs about the individual loudness levels of sounds and their ordinal rankings are generally true, but our beliefs about their ratio relations are false. So, in particular, Soft's report is just false.

But this response is unsatisfactory for three reasons. First, there is empirical evidence in support of the truth of ratio reports concerning apparent loudness, as I already noted. Further, if ratio judgments concerning loudness makes no more sense than ratio judgments concerning level of beauty (say), then it is a wonder that subjects fairly consistently make such judgments at all.

Second, if objectivism is true concerning one range of sensible qualities (e. g. audible qualities), it is presumably true for other ranges (e. g. smell qualities). So, although O'Callaghan does not consider other cases, the objectivist needs a response in every case. Now, even if we accept an error theory about Soft's sophisticated ratio report, it is much harder to accept an error theory about simple judgments of relative resemblance, such as Sniff's report that one smell quality is more like a second than a third. Indeed, such reports about resemblances among general qualities (as opposed to particular items in the environment), such as *blue is more like purple than yellow*, are often thought to be certain *a priori*.

But if this does not convince you, let me make a third point against the error theory. We not only make structure judgments about the sensible qualities apparently possessed by items in the external world; we

can also make structure judgments about our own experiences. For instance, consider Sniff's introspective judgment that the citrus smell of limonene *seems* more like the citrus smell of citral than the minty smell of menthol, in other words, that his smell *experience* of limonene is more like his smell *experience* of citral than his smell *experience* of limonene. What is the right account of phenomenal reports like this?

Even though I reject *externalist* intentionalism, I think that the basic intentionalist approach to phenomenology is plausible. On intentionalism, all phenomenal facts derive from facts about the properties presented in experience. (Maybe there are *some* exceptions – for instance, involving visual blur.) So, presumably, Sniff's consecutive smell experiences fall into the relevant phenomenal resemblance-order just in case the successively presented smell qualities do. The phenomenal structure of experiences is inherited from the structure of the qualities successively presented in those experiences. Call this the *inheritance claim*. On intentionalism, there appears to be no other option. The inheritance claim is actually intuitively plausible in general, independently of any theory. As Alex Byrne (2003, 645) says: "Why is the experience as of a teal object [phenomenally] similar to the experience as of a turquoise object? Because teal is similar to turquoise".

But, on the error theory of the structure of the sensible qualities, the smell qualities that Sniff is presented with are objective chemical properties that *do not* stand in the relevant resemblance-order. Hence, given the inheritance claim, the error theory *spreads* to Sniff's introspective judgment that the citrus smell of limonene *seems* more like the citrus smell of citral than the minty smell of menthol. The error theory implies that Sniff is even wrong in thinking that that his smell *experience* of limonene is more like his experience of citral than his smell experience of limonene. By similar reasoning, the error theory contemplated by O'Callaghan entails that even our introspective judgments about the *apparent* magnitude relations (ratios, equal differences) among sounds are false. But this is hard to accept. It is especially hard to accept an error theory when it comes to Sniff's simple introspective judgment about resemblance-order.

Complex respondent-independent theory. So we need a theory on which our structure judgments come out true. The externalists have two options: a *response-independent theory* of qualitative structure, which does not appeal to our internal neural or behavioral responses to external properties in any way; or a *response-dependent theory* of qualita-

tive structure, which does somehow appeal to our internal neural or behavioral responses.⁹

Now, given good internal correlation and bad external correlation, our structure judgments “match” our internal neural responses better than the external conditions that prompt them. So a response-dependent theory of qualitative structure seems like an obvious choice; and a response-independent theory seems hopeless. Nevertheless, I will start with the response-independent theory of qualitative structure because such a theory fits best with my targets here, tracking intentionalism and objectivism. The motivation behind such views is that sensory character seems wholly “out there” in the response-independent world. And Michael Tye (2000, 163) has actually attempted a response-independent theory of simple color structure judgments like “purple is reddish and bluish”. Although he does not generalize to more complex judgments (e. g. judgments of relative resemblance) or other ranges of sensible qualities, his discussion hints at a general strategy for devising such accounts.¹⁰

The strategy is simple. Structure judgments are correlated not just with internal properties and relations but also with external ones. In some cases, the external correlates are revealed by psychophysics. So corresponding to every response-dependent account there will be a response-independent account. Given bad external correlation and good internal correlation, the external correlates of structure judgments will be more complex and unnatural than the internal correlates. Nevertheless, according to Tye, the subject matter of our structure judgments are the external, response-independent correlates.

So, for instance, psychophysics has revealed that, as a rough rule of thumb, for simple tones with the same frequencies, we judge that loudness has doubled just in case physical intensity (related to amplitude) increases tenfold. So, on a grossly oversimplified response-independent theory, loudness levels are just physical intensities and when we use ‘twice greater’ than in relation to loudness it takes on a new semantic value: it comes to mean *ten times greater*, although we are

⁹ Since tracking intentionalism is meant to be a reductive theory of consciousness, I will be ignoring “primitivism” about the sensible qualities; so the only response-independent theories I will consider will be *reductive* response-independent theories.

¹⁰ Tye (2000) and also Byrne and Hilbert (2003) defend an interesting *hue-magnitude account* of color structure. But even if that account is correct in the case of color (I think there are several problems with it), the same account obviously does not apply to the qualitative structure of smells or audible qualities. So objectivists would need some other accounts here; as we shall, it is very hard to see what accounts they might provide.

semantically blind to this! In this way, an objectivist about loudness like O'Callaghan might say that our "doubling" judgments in relation to loudness come out true. Now, in fact, whether one sound appears twice louder than another also depends on a complex variety of other factors, including frequency, duration, and critical bands. Moreover, the relationship between loudness and intensity changes for levels below 40 dB SPL. So, the proponent of a response-independent theory of auditory structure would have to say that loudness levels are extremely complex physical properties and that the semantic value of 'is twice greater than' in auditory contexts is some horrendously complex, disjunctive relation that cannot be easily defined. In order to ascertain this complex relation, we look at our responses. But the relation itself is response-independent.

Likewise, the proponent of a response-independent theory of qualitative structure might claim that pains are complex bodily disturbances and that in contexts where we use 'twice greater than' in relation to bodily disturbances *D2* and *D1* we are referring to the relation which is the disjunctive psychophysical correlate of our pain-doubling judgments: *D2* and *D1* are noxious temperatures with such-and such spatial extents and *D2* is 33% greater than *D1*, or *D2* and *D1* are electric shocks applied to the skin with so-and-so spatial extents and *D2* is 25% greater than *D1*, or . . .

It would be even more difficult to develop a wholly response-independent theory of qualitative structure in the domain of smell. Such a theory would require that there is a single response-independent relation *R* such that we (or smell experts) judge *smell S1 is overall more like S2 more than S3* when and only when the corresponding molecular types *C1*, *C2* and *C3* stand in relation *R*; and that this relation *R* is the semantic value of 'x is overall more like y than z' in our reports of smell similarity of the form *smell S1 is overall more like S2 more than S3*. For, on a purely response-independent theory, only then will such reports come out generally *true*. But the psychophysics of smell is particularly messy. There simply is no single response-independent relation *R* that fills the bill.¹¹

¹¹ Haddad et al. (2008) and Turin (2002) attempt to relate chemical similarity to qualitative similarity. But they are very far from establishing what a response-independent theory of smell similarity would require: that there is a single response-independent relation *R* such that we (or smell experts) judge *smell S1 is overall more like S2 more than S3* when and only when the corresponding molecular types *C1*, *C2* and *C3* stand in relation *R*. First, even when it comes to simple monomolecular odors, their methods are open to counterexamples and do not come close to explaining all of the variance. Second, they do not take into the account the effects of concentration on quality,

Of course, those who favor a wholly response-independent theory of smell similarity might say that, when we use the predicate ‘ x is overall more like y than z ’ in connection with smells, the predicate comes to express response-independent relation that only be defined in terms of a “big list”: a list of all the chemical structures $C1$, $C2$ and $C2$ corresponding to the smells $S1$, $S2$ and $S3$ such that we judge that $S1$ is more like $S2$ more than $S3$. This would guarantee that our paradigmatic smell similarity judgments come out true, but it would be it would be totally implausible from a semantic point of view.

Now that we have attempted to elaborate the response-independent theory of qualitative structure, we can see that it is quite hopeless. There are many problems. Let me just mention the simplest one, which also goes deepest. I call it the *metasemantic objection*. The response-independent theory is in part a theory of the content or truth-conditions of our talk about the structural features of sensible qualities. It says that such talk about sensible qualities and their relations is about very complicated response-independent properties and their response-independent relations. But no *metasemantic theory* of how language hooks up to the world would support this view. Any metasemantic theory would instead support a response-dependent theory.

To see this, consider an example. As we saw, on the response-independent theory, when we speak about sounds, the predicate ‘ x is twice more intense than y ’ comes to express a horrendously complex response-independent relation I . This relation must be extremely complex, because of response compression as well as the dependence of loudness on multiple objective factors. Because the relation I is response-independent, on this theory our ratio judgments about sound-intensity have wholly response-independent truth-conditions. By contrast, on a simple response-dependent theory, when we speak about sounds, the predicate ‘ x is twice more intense than y ’ expresses a quite

which can be extreme (Malnic *et al.* 1999). Third, their methods only apply to simple monomolecular odors. They do not apply to natural odor objects, such as the odors of foods and plants, which are typically mixtures containing tens to hundred of monomolecular components, and whose smells are not at all a function of the smells of their components. Could it be that some future canonical physical description of molecular types will make it evident that there is a relation R among molecular types that perfectly tracks our smell similarity judgments? (In an interesting discussion, Davies (forthcoming) brings up a similar question regarding reflectance properties and judgments of color resemblance.) All the evidence suggests that the answer is ‘No’. Moreover, we must remember that the proponent of a response-independent theory of qualitative structure would need to show how to accommodate our structure judgments about all types of sensible qualities (loudness levels, pitches, tastes, etc.) in purely response-independent terms. This seems impossible.

different, response-dependent relation *D*. In one version, sounds are external physical events, and *D* is a relation definable along these lines: *x and y normally cause a doubling of total neural response in the relevant auditory channel*. Now, evidently, *D* is much more simple or “natural” than *I*. Further, *D* plays more of a role than *I* in causally explaining our judgments about loudness doubling. Now nearly all physicalist metasemantic theories of content or truth-conditions of our language appeal to causation or naturalness or both. Therefore, given bad external correlation and good internal correlation, any physicalist is more or less *forced* to accept some response-dependent theory of the truth-conditions (contents) of our judgments of loudness scaling. Likewise for judgments about pain intensity and judgments about relative resemblances among smells.

Respondent-dependent theory. So, could tracking intentionalists (Tye, Dretske) and other objectivists about sensible qualities (Hilbert, Byrne, O’Callaghan, Batty) block the structure argument by accepting some kind of response-dependent theory of qualitative structure?

To begin with objectivists certainly cannot accept *some* response-dependent theories. It is part of their view that individual sensible qualities (audible qualities, taste qualities) are *objective, response-independent* properties of external items. So they obviously cannot hold that the sensible qualities are internal neural properties projected onto the external world (projectivism) or dispositions to produce internal neural responses (dispositionalism).

But, on the face of it, tracking intentionalists and other objectivists about sensible qualities apparently could hold that, while sensible qualities are objective physical properties, their *structural relations* are to be explained in terms of the responses they produce in us. On this mixed view, loudness levels are objective physical properties, but when we judge that one is roughly twice greater than another our judgment is true just in case the one normally causes a total neural response roughly twice greater than that normally caused by the other. Likewise, felt pains are objective bodily disturbances, but their sensory intensity is explained in terms of the neural response they typically cause. And smells are objective chemical phenomena, but when we judge that they resemble to a certain degree our judgment is true just in case they normally produce similar ensemble activations in the PPC (the only known physical correlate of smell similarity). If you like, the smells are objective phenomena but they are similar or different *in respect of* their effects on us. One might think that, in this way, even tracking inten-

tionalists and other objectivists about sensible qualities can accommodate the structure judgments of Sniff, Soft and Mild.

However, even this proposal is unsatisfactory. To begin with, my main target here is tracking intentionalism. The response-dependent response to the structure argument is not available to tracking intentionalists for the simple reason that it is *inconsistent with* tracking intentionalism. On the response-dependent theory, the individuals in coincidental variation cases (Yuck and Yum, Sniff and Snort, Mild and Severe, Soft and Loud) have *phenomenally different* experiences of the same objective properties because those properties typically produce different neural responses in them. In particular, their experiences *exhibit different phenomenal structure*. But, as we have seen, tracking intentionalists must apparently say that they have *phenomenally identical* experiences, because the track and hence represent exactly the same external conditions. Only the hopeless *response-independent* theory of qualitative structure appears consistent with the radical externalist theory of phenomenal character promoted by tracking intentionalists.

Now, you might think that, even if tracking intentionalists cannot accept the response-dependent response to the structure argument, others who favor objectivism about sensible qualities (O'Callaghan, Batty) might accept it, if they simply reject tracking intentionalism. But there are serious problems for this response, even if we ignore tracking intentionalism.

First of all, as I have mentioned, even if tracking intentionalism fails, I do believe that there are powerful reasons to accept some version of intentionalism about experience. But the response-dependent theory of qualitative structure is hard to square with *any* version of intentionalism, not only tracking intentionalism. For instance, on the response-dependent theory, Sniff and Snort's smell experiences differ in phenomenal structure, because their neural (PPC) responses differ. But they apparently represent the same individual objective (perhaps disjunctive) chemical types, because they bear the same naturalistic relations to them. According to the objectivist who accepts the response-dependent theory of qualitative structure, what then is the representational difference between their experiences that constitutes the phenomenal difference? It is not enough to say that their different neural responses constitute the phenomenal difference; given intentionalism, this neural difference must be accompanied by a representational difference (for an argument, see the discussion of "quasi-intentionalism" in §5).

Perhaps the objectivist who accepts the response-dependent theory of qualitative structure will reply that Sniff and Snort's experiences do not only represent the objective (chemical) properties of odors; they also represent different contents of the form *the odor clouds have objective properties that cause in me PPC states that resemble to degree D*. Their olfactory systems are "self-centered", because they represent, not only the objective properties of things, but also response-dependent properties concerning the effects of those things on those very systems. It is because Sniff and Snort's experiences differ in these "self-centered" contents that their experiences differ in phenomenal character.

But how could Sniff and Snort's experiences manage to represent such bizarre, complex contents? On the face of it, there is no naturalistic theory of intentionality compatible with this view. (For more on this point, see my response to Kriegel's related view in §5.) Further, in any case, since Sniff and Snort smell the odors consecutively, at no time could their experiences represent such contents involving multiple odors.

There is another decisive problem with the combination of an objectivist theory of individual sensible qualities and a response-dependent theory of their qualitative structure. Evidently, if one loudness level is twice greater than another, then this is an essential feature of the loudness levels. Compare: if one length is twice greater than another, this is an *essential* feature of the lengths. Likewise, intuitively, if two smell qualities resemble to degree *D*, then this is an *essential* feature of them, one they possess in any situation in which they exist. But, on the combination of an objectivist theory of individual sensible qualities and a response-dependent theory of their qualitative structure, these intuitions are false. For, on objectivist theory of individual sensible qualities, these sensible qualities are objective physical properties. Further, those very properties might of course have normally produced quite different neural responses than they in fact do. (This, indeed, is what happens in coincidental variation cases.) On a response-dependent theory of qualitative structure, this means that those very sensible qualities might have had quite different qualitative structure than they in fact do.¹²

This concludes my defense of the structure argument. While others have discussed similar arguments concerning color, I have developed a

¹² For other problems with the combination of a response-independent (objectivist) theory of sensible qualities and a response-dependent theory of their qualitative structure, see Pautz (2006b).

new version of the structure argument concerning other sensible qualities. And I have shown why the usual responses are unsatisfactory. The conclusion I draw is that tracking intentionalists and objectivists about sensible qualities can provide no satisfactory account of the truth of our ordinary qualitative structure judgments, such as those of Sniff, Soft and Mild. This is perhaps the best way of stating the structure argument against these views.

5. No Refuge for Externalist Intentionalists

I have taken tracking intentionalism as my stalking horse. Some phenomenal externalists (Hill, Lycan, Tye, others) have suggested that the kind of tracking intentionalism I have focused on is too simple and that more complex versions of externalist intentionalism avoid my arguments. Therefore I conclude by rebutting their responses to my arguments.

Valuational properties. Recently, Michael Tye and Brian Cutter (2011) have claimed tracking intentionalism can handle cases like that of Mild and Severe. Recall that in response to increasing noxious temperatures Severe's S1 firing rates (known in our own case to be linearly related to sensory intensity) increase much more rapidly than Mild's. He also produces higher pain ratings of sensory intensity the VAS scale. To handle this sort of case, Tye and Cutter retain the claim that it is the representation of objective features of the disturbance (size, intensity) that determines *sensory intensity*, even though studies show that there is at best a complex non-linear relationship here. Since they agree that tracking intentionalism entails that Mild and Severe's experiences of increasing noxious temperatures represent bodily disturbances of exactly the same objective types, they must say that their pain experiences of these noxious temperatures are exactly alike in *sensory intensity*, despite the strong neural and behavioral evidence to the contrary (this was confirmed in correspondence).

However Tye and Cutter would allow that there is *some* phenomenological difference between them. In cases like this, they claim that, while Mild and Severe's experiences represent the same bodily disturbance types, they also represent different *valuational properties*. Their experiences have "layered content". The same disturbance is bad-for-Mild-to-degree-y and bad-for-Severe-to-degree-y, where y is greater than x. They explain badness in terms of aptness for harm but provide

no naturalistic explanation of across-species comparisons of *degrees* of badness. This is a serious lacuna in their account. In any case, their view is that Mild's experience represents the disturbance as having the first valuational property while Severe's experience represents the disturbance as having the second one, because their pain systems actually track these different valuational properties. Further, in consequence of this representational difference, their pains differ in *affective phenomenology* or *unpleasantness*, even if they are exactly alike in *sensory intensity*. The (well-documented) difference between *sensory intensity* and *unpleasantness* is subtle. Roughly, Tye and Cutter's idea is that Mild and Severe have pains of exactly the same sensory intensity, but a pain of that intensity *bothers* Severe more than it does Mild. This is supposed to explain the fine-grained behavioral differences between them. This is their answer to my "internal-dependence" argument about pain.

In response to my point that there is "bad external correlation" in the case of pain, Christopher Hill (2012, 137) has independently invoked a notion somewhat similar to the notion of badness level. He suggests that it makes sense to claim that the *threat level* of one disturbance is *twice greater than* another. There is a lacuna in his account, because he does not say what exactly what naturalistic facts constitute a doubling of threat. In any case, the idea is that, if Mild (who, recall, is an actual individual) experiences a small increase in temperature from 47°C to 50°C, the threat level might actually double. This might accommodate the truth of his report that his pain intensity doubled, thus answering my "structure argument" about pain. Hill might also say that Mild and Severe represent the same bodily disturbances as having different threat levels. He might say that this means that their pains differ in phenomenology. He might even say that this means that they consequently differ in *sensory intensity*, not merely affective phenomenology. In that case, his view would be somewhat different from that proposed by Cutter and Tye, who claim that there is only an affective difference.

As Tye and Cutter note, *taste experiences* might also represent valuational properties. So the externalist intentionalist might claim that, while the fine-grained sensory character (sweet, salty, bitter, sour, umami) of taste experience is determined by the representation of chemical types (sugar, salt, acid, etc.), the affective character of taste (good or bad) is determined by the representation of valuational properties (good or bad). Now recall the case of Yuck and Yum. The externalist might claim that, while their experiences represent the same

chemical property *C* of the berries, Yuck's experience somehow represents the berries as bad for him (to some degree) and Yum's experience somehow represents them as good or maybe "edible" for him (to some degree). It's not just Yuck and Yum *think* these things; the idea is that their *experiences* comment on the nutritional value of the berries! In consequence, the tracking intentionalism might claim that, while their taste experiences are identical in sensory character, they differ on the affective dimension. For instance, maybe both have the same strongly bitter experience, but strangely Yum really *likes* it.

The view being proposed, then, is that *some* of our experiences represent valuational properties, in addition to objective properties. Call this the *valuational view*. And this is supposed to help answer my arguments.¹³

My Reply. In fact the valuational view does not answer my arguments. First, and most importantly, it does not answer the internal-dependence argument when it comes to *Sniff and Snort* or *Soft and Loud*. Here the external stimuli (odor clouds, sounds) do not differ in "valuational properties". For instance, the same sounds are not "good" for Soft and "bad" for Loud. So here the valuational gambit does not get off the ground. It also does not provide an answer to the structure argument. Even if externalists like Hill can use threat levels to accommodate Mild's structure judgments about pain levels, they obviously cannot use them to accommodate Soft's judgment about *loudness* levels, or Sniff's smell resemblance judgment. Since they defend a general theory of phenomenal consciousness, Tye and Cutter and Hill need to say something about these other apparent counterexamples to their theories.

In fact, the valuational view does not help answer my argument about Mild and Severe. Tye and Cutter discuss a version of the Mild and Severe case in which the same disturbance is more of a threat to

¹³ It might be thought that I should also consider here imperative intentionalism about pain defended by Klein (2007) and Martinez (2011), since externalist intentionalists might naturally use that theory to answer my arguments. But, for two reasons, I will not consider that view separately here. First, the main points I will make about the valuational account apply equally to the imperative account: for instance, it cannot be applied to all of my cases, so it would not afford a general solution to the problems I raise here for externalist intentionalism. Second, elsewhere (2010, note 36) I suggest that it faces especially serious problems. Klein and Martinez (forthcoming) suggest that they can handle *one* of the problems I raise there about degrees of pain (a problem repeated by Cutter and Tye (2011)), but even if they are right I believe that the other problems I list are enough to show that imperative intentionalism about pain is difficult to defend.

Severe than to Mild. But, in my present version of the case, I stipulated that this is not true. The noxious stimuli and their aptness to harm are held constant, while there is variation in S1 activity and VAS pain ratings of sensory intensity. Therefore, even if we grant that “degrees of badness” make sense and can be represented by experience (which I question below), the tracking intentionalist is stuck with the mistaken verdict that Mild and Severe’s experiences are in every respect exactly alike in phenomenal character.

The proposal of Tye and Cutter also fails when it comes to the kind of case they actually discuss, where the same disturbance is more of a threat to Severe than to Mild. In this kind of case, their proposal at best entails that their Mild and Severe’s pains only differ at the affective level, and are exactly the same in sensory intensity. But Tye and Cutter ignore a feature of the case that I have stressed here and elsewhere (Pautz 2010). My argument and the empirical evidence I adduce exclusively concern *sensory* intensity. Pain intensity at the sensory level is linearly related to *S1 firing rates* and is at best related in a complex non-linear fashion to objective features of our bodily disturbances. As for the affective dimension of pain, impressive fMRI studies by Hofbauer and Rainville and others show that it is coded by activity in the anterior cingulate cortex (ACC), as I noted previously (§2). Now I stipulated that Mild and Severe differ in *S1 activity* (not just ACC activity), and their responses on the visual analogue scale *for sensory intensity*. Given these points, the only plausible verdict concerning the case is that their pains differ in *sensory* intensity.

So, both versions of the case of Mild and Severe are in fact counterexamples to tracking intentionalism, even if we take valuational properties into account. The case of Yuck and Yum, too, is a counterexample to tracking intentionalism. As we saw, on the valuational view, their experiences are *identical in sensory character* (because they represent the same chemical type), and only differ in the affective dimension of taste. For instance, maybe both have the same bitter experience, but strangely Yum really *likes* it. This verdict simply does not fit the facts. It is clear that ensemble activation patterns code the sensory character of taste experience (sweet, bitter, etc.), not the affective character. This is shown by the fact that tastes that differ in sensory character but agree in valance (e. g. bitter and sour tastes) are realized by different ensemble activation states the brain. Since Yuck and Yum’s ensemble activation states occupy different positions in neural taste space, and since their fine-grained sorting and other behaviors differ, it

is totally implausible that they both have (say) a very bitter experience of the berries. Instead, their experiences differ in sensory character. So this case shows that tracking intentionalism fails, even if we take valuational properties into account.¹⁴

Finally, Hill's interesting threat-level proposal for answering my structure argument about pain faces some problems. For one thing, Hill does not say what naturalistic facts might ground a *doubling* in threat-level. When Mild goes from feeling 47°C to feeling 50°C, the probability that he will die does not double. What then could it mean to say that the threat-level doubles? And, again, Hill's proposal evidently does not apply to Soft's judgment about a doubling of perceived loudness. Further, in both of these cases, there is a natural alternative proposal: the appeal to *firing rate*, which is well-defined (unlike "threat-level") and which is well correlated with sensory intensity. This seems to be what is driving our judgments about "intensity". How then could a

¹⁴ I have another worry about Tye and Cutter's specific version of the valuational gambit. Maybe Millikan's teleological (1989) theory of representation is compatible with their claim that our experiences represent properties like *being bad* or *being poisonous*. (In fact, a problem with her theory might be that it entails that our experiences *only* represent such properties.) By contrast, Tye's own theory of representation appears incompatible with that claim. On his theory, in order for a state to represent a valuational property like *being bad* or *being poisonous*, its tokenings must be *explained by* the instantiation of that property under normal conditions. Despite Tye and Cutter's (2011) interesting efforts to show that this condition is fulfilled, I still have doubts. To see why, notice that, whenever Yuck tastes the berries and his neural representation of the berries is tokened, the following counterfactual is true: if the berries were not poisonous (bad) to Yuck's species (if e. g. we gave them a pill that prevents the action of the poison, or if we somehow removed the poisonous part of the berries), that neural representation still would have been tokened (because the chemical properties of the berries still would have impinged in the same way on the taste system). (See also Pautz 2010, note 15.) The truth of this counterfactual shows that the tokening of the neural representation on particular occasions is not *explained by* the poisonousness, or badness, of the berries; rather, it is only explained by the response-independent chemical property of the berries. A similar counterfactual-based argument would show that the tokenings of pain representations in Mild and Severe is never explained by the "badness" (or badness-to-degree-x) of pain stimuli. So, on Tye's theory of representation, experiences cannot represent these properties. Tye and Cutter (2011, p. 100-101) argue that the explanatory condition on representation is fulfilled because the badness of the stimuli provides a historical explanation for why *these species were designed by natural selection to token such states that cause withdraw*. This may be true. But, contrary to what they seem to think, it does not follow (and is in fact not true, in view of my counterfactual argument) that *particular tokenings* of these states at the present time are ever explained by the badness of the stimuli, which is what their theory requires. (Compare: a fire alarm might have been designed to ring because fire is dangerous, but particular occurrences of the ring in the present are explained by the presence of smoke, not danger.) Further, even if this problem can be overcome, since Tye and Cutter only attempt to explain how experience represents course-grained evaluative properties like *being bad*, they would still need to explain how experience represents one fine-grained degree (*being-bad-to-degree-x*) rather than another (which would require providing the naturalistic grounds for across-species comparisons of degrees of badness).

naturalist avoid the conclusion that our talk of “intensity” in these domains somehow *refers to* overall firing rate or something along these lines?

Pythagoreanism about sensible magnitudes. Casey O’Callaghan (in discussion) suggested this view to me in the case of loudness, without wholeheartedly endorsing it. (On one way of elaborating the hue-magnitude theory of color defended by Byrne and Hilbert (2003) and Tye (2000), it is a version of what I will call “Pythagoreanism”.) Since the view is somewhat difficult, let’s start by seeing how it applies to an extremely simple and fanciful coincidental variation case. Then we will turn to my cases.

Suppose that A and B are two devices. In each device, there is a cylinder of fluid. The devices somehow respond to the lengths of objects in the environment and encode these lengths in terms of the height of the fluid. However, in each case, there is an expansive non-linearity. There is “bad external correlation”. In the case of A, a line of n cm results in A’s inner fluid rising to n^2 cm. In the case of B, the “response curve” is steeper: a line of n cm results in A’s inner fluid rising to n^3 cm. Let us also pretend that A and B have sensations that vary in intensity (rather like auditory experiences or pains), where the intensity is linearly related to inner the fluid level. There is “good internal correlation”. (If you like, pretend that they are alien sensations of a sort that we do not have.) Thus, when they both respond to (“track”) a line of 2 cm and then a line of 3 cm, A’s sensation roughly *doubles* in intensity (because his inner fluid level goes from 4 cm to 9 cm), while B’s sensation *more than doubles* in intensity (because his inner fluid level goes from 8 cm to 27 cm). In addition, A “judges” that the intensity has doubled, while B “judges” that it has more than doubled.

This provides a schematic illustration of my internal-dependent argument and my structure argument. How could the “tracking intentionalist” or “objectivist” about sensible qualities accommodate the verdict that A and B have sensations of different intensities, and how might he accommodate the truth of their structure judgments? After all, they track the very same objective lengths; and the objective length does not really *double* (or more than double) when it goes from 2 cm to 3 cm. Given bad external correlation and good internal correlation, how could their sensations be mere representations of external lengths? Isn’t it more natural to take an internalist approach, on which their sensations are identical with, or supervene on, inner fluid levels?

The Pythagorean response is meant to save tracking intentionalism and objectivism as follows. Let us say that a line has an *x-value of x* just in case it has a length of n cm and $x = n^2$. And let us say that a line has a *y-value of y* just in case it has a length of n cm and $x = n^3$. Suppose A and B are presented with a line of 2 cm. The Pythagorean holds that the line instantiates the following three properties: the physical length l (which does not involve number in any way, although we assign numbers to it), the property *having an x-value of 4* and the property *having y-value of 8*. He accepts an extremely fine-grained view of properties, on which these properties are distinct, even though they are necessarily co-extensive. The idea is that the second two properties are relations to numbers; since they are relations to different numbers, and involve different functions, we should count them as non-identical, even if they are necessarily co-extensive with each other. (Likewise, he would even hold that the property of bearing the length-in-centimeters relation to the number 100 is distinct from, but necessarily co-extensive with, the property of being the length-in-meters relation to the number 1.) Further, according to the Pythagorean, on being presented with a line of 2 cm, A represents (“perceives”) the property *having an x-value of 4* and B represents (“perceives”) the distinct, but necessarily co-extensive, property *having y-value of 8*. This representational difference constitutes the difference in intensity between their sensations, according to him. This answers the “internal-dependence argument”. As for the structure argument, notice that x-values perfectly match A’s structure judgments. Thus, when the line goes from 2 cm to 3 cm, and A judges that the intensity has doubled, the represented x-value doubles (from 4 to 9). So, on this view, his structure judgments come out true. This is not a subjectivist or response-dependent view. A’s structure judgments are not about his inner fluid levels. Rather, they are about the represented x-value properties. And these properties are response-independent properties of things (even if they match A’s inner fluid levels). For instance, even if A were not around, a line of 2 cm would have an x-value of 4.

Of course, I call this the “Pythagorean view” because it holds that experience represents properties defined in numerical terms. I think that this is an essential feature of the view. Suppose someone just said that, in addition to lengths in cm, the line has two other families of properties, and that, while the lengths do not match A and B’s structure judgments, these other properties do. We would be mystified. To

make the view comprehensible, he must define x-values and y-values, and explain that he individuates properties extremely finely.

Now that we have a grip on the view, we can see how it might apply to more complex cases, for instance, the case of Soft and Loud. Of course, that case is somewhat analogous to the case of A and B. The main difference is that, while A and B's experiences track non-disjunctive properties (corresponding to individual lengths), Soft and Loud's experiences track extremely complex, disjunctive properties involving amplitude, frequency, critical bands, and so on. The Pythagorean about sensible magnitudes holds that there is a complex function f from these parameters onto numbers that reflects Soft's psychophysical judgments of loudness level. We might call this the S-loudness of a tone. There is a different complex function g from these parameters onto numbers that reflects Loud's (different) psychophysical judgments of loudness level. We might call this the L-loudness of a tone. (Those who work in psychophysics actually have devised scales that reflect our psychophysical judgments, such as the Bark scale and the mel scale.) According to the Pythagorean, the loudness levels that Soft perceives are S-loudness levels, while the loudness levels that Loud perceives are L-loudness levels. So for instance if they hear the same tone, the tone might have both the property *having an S-loudness of 10* and the property *having an L-loudness of 20*. These properties will be necessarily co-extensive, because the very same combinations of physical parameters that yield an S-loudness of 10 yield an L-loudness of 20. But the thought is that they are nevertheless distinct, and that Soft some perceives (or "represents") the first one while Loud perceives the second one. So Pythagoreanism about loudness levels answers my internal-dependence argument. Further, since these properties match their ratio judgments, it also answers my structure argument. Maybe a similar view could be applied to the case of Mild and Severe involving pain intensity (Pautz 2010). This counts as a sophisticated version of tracking intentionalism. The idea is that sensory magnitudes are objective (if complex) properties and sensory intensity is determined by what sensory magnitudes we track.

My reply. The main problem with Pythagoreanism is that it does not provide a general response to my arguments. The mathematical treatment of sensory qualities certainly does not apply to smell qualities or taste qualities, for instance. Therefore the Pythagorean response does not help with my arguments concerning smell and taste. Those argu-

ments are enough to undermine externalist intentionalism and objectivism about sensible qualities more generally.

In fact, for several reasons, Pythagoreanism even fails in the case of loudness. (i) While we can in a rough and ready way *represent* loudness levels in terms of numbers, any *definition* of loudness levels in mathematical terms is totally implausible. There is conventionality and vagueness involved. (ii) To appreciate my next problem, return to the simple case of A and B. The Pythagorean view requires that *having an x-value of 4* is distinct from *having y-value of 8*, even though each is necessarily coextensive with the same length (having a length of 2 cm). Against this, intuitively, what we have here are just two different descriptions of a single length property. (iii) Even if we grant that necessarily coextensive properties can be distinct, the view certainly fails. To appreciate the problem, suppose again that A and B are presented with a 2 cm line. The Pythagorean holds that A's experience represents the property *having an x-value of 4* while B's experience represents the property *having y-value of 8*. But, given that these properties are necessarily co-extensive, what makes it the case that A's experience represents *having an x-value of 4* but not *having a y-value of 8*, while B's experience represents *having a y-value of 8* but not *having an x-value of 4*? The proponent of this view needs a *general theory* of the sensory representation relation, the relation *x sensorily represents property y*, which explains how this might be so. But those theories always appeal to relations like tracking or indication, which cannot distinguish between necessarily co-extensive properties. It would not be enough for the Pythagorean to respond by rejecting this kind of theory. He would need to at least gesture at an alternative general theory of the sensory representation relation that answers the problem. (iii) Supposing we can make sense of Pythagorean properties, they are evidently extremely disjunctive and unnatural. It would seem that firing rates in the auditory channel are much more natural. Since they seem well correlated with auditory intensity, and since they seem to be what causally explains judgments about 'loudness', the naturalistic *must* claim that 'loudness' refers to a property involving firing rates in the auditory channel. By similar reasoning, he *must* say that 'pain intensity' refers to a property involving firing rates in S1 or other regions of the pain-matrix. There is no other reasonable view for the naturalist. This naturally leads to "neural projectivism".

Neural projectivism. This interesting view was suggested to me by Christopher Hill as a possible view of pain intensity. (Hill did not outright defend the view, but only suggested it as a possibility.) On this view, when you have a throbbing pain in your foot, for instance, there is in your brain an atomic representation R of intensity. But this neural representation R does *not* represent an actual peripheral state, such as stimulus intensity or size. Bad external correlation means that pain intensity is not directly proportional to any such peripheral stimulus features. Rather, R represents the “intensity” of the *internal nociceptive neural signals* set up by the stimulus, which are known to be more proportional to pain intensity (“good internal correlation”). Now, the “intensity” of nociceptive neural signals is just a matter of firing rate. So the idea is that R represents a firing rate property of the form *exhibiting firing rate N* . (Of course, even if pain intensity is firing rate, our pain experience doesn’t reveal that it is a matter of firing rate, just as a perception of water does not reveal that it is H_2O .) In one version of neural projectivism, R represents the firing rates of nociceptive neurons in the spinal cord. In another version, R represents the (possibly different) firing rates of nociceptive neurons in the cortex itself (e. g. S1), which are known to be especially well correlated with pain intensity. (I will suggest below that this is the best version.) Thus R is a kind of “self-monitoring” representation.

Now for the “projective” element of the view. The neural projectivist further claims that, in the area of the brain responsible for pain, there are other atomic representations, R' , R'' , . . . of other features. By contrast to R , these atomic representations *do* represent *peripheral* properties, like *throbbing* and *being located in the foot*. According to the neural projectivist, when you have a pain, the atomic representation R of *exhibiting firing rate N* is “combined with” these other representations R' , R'' , . . . The result is a complex representation with the content *there is something with the properties throbbing, being located in the foot, and exhibiting firing rate N* . This content is false: while there is disturbance with the properties *throbbing* and *being located in the foot*, it does not have the property *exhibiting firing rate N* . In that sense, your pain experience projects a property in fact possessed by the central nervous system onto a external bodily region. Maybe this could be thought of as a kind of “binding error”, because it involves binding a property that is possessed by one entity (the nervous system) together with properties that are possessed by another entity.

That, then, is neural projectivism. *If* neural projectivism about pain is workable, it answers my arguments about pain. On neural projectivism, Mild and Severe have pains of different intensities in response to the same noxious temperatures, because they “project” different firing rate properties (in fact possessed by their central nervous system) onto the external bodily regions to which the temperatures are applied. This answers my internal-dependence argument. On neural projectivism, when in the actual world Mild judges that his pain has roughly doubled in intensity between 47°C and 49.8°C, his judgment is actually about the neural firing rate projected onto the relevant bodily region. If there really is a linear correlation between neural firing rate and pain intensity, the firing rate literally doubles. So, Mild’s ratio judgment is literally true. This answers my structure argument.

Given that there is “good internal correlation” and “bad external correlation” in all sense-modalities, if neural projectivism is right about pains it must be right about other experiences.¹⁵ The idea would be that loudness is neural firing rate in the auditory channel; smell and taste qualities are ensemble activation patterns; color qualities are neural properties of the chromatic channels in the brain; and so on. And all of these sensible qualities are somehow projected onto external events and items, according to the uniform neural projectivist. They are bound together with external locations, shapes, and so on. Such a view might answer my arguments about Yuck and Yum, Sniff and Snort and Soft and Loud. Of course, this would just be a naturalistic version of the traditional Galilean view that locates the sensible qualities in the head. It would be a projectivist version of the “stinking brain” theory (when color is at issue, “colored brain” theory).

Since externalist intentionalists were trying to avoid exactly this view (recall the quote from Armstrong at the start of the paper), one might wonder whether we should count it as a version of externalist intentionalism. I think it does deserve the name. On neural projectivism, many representations really do accurately represent external properties, like location and shape, presumably by way of some kind of

¹⁵ For instance, our perceptions of the four perceptually prominent elemental or “unique hues” cannot be explained merely in terms of the reflectance properties of external objects. Most researchers assume that it has a cortical basis but for many years that remained elusive. However, a recent breakthrough study by Horwitz and Haas (2012) seems to have gone some way towards uncovering the cortical basis of such perceptions. So there is reason to think that the consistent neural projectivist about pain intensity would have to be a projectivist about the qualitative dimensions of color as well.

tracking or indication relation; and the representation of such properties plays an important role in configuring phenomenology.

My reply. Before I evaluate neural projectivism, let me say that I think the best version of this view holds that the projected neural properties are properties that in fact belong to *cortical* neural assemblies rather than more peripheral neural signals. The reason is that this version of the view is needed to handle the full range of hypothetical “coincidental variation cases”. Coghill *et al.* (2003, 8542) report that in some actual cases there is reason to think that “a large portion of the variability of interindividual differences in both the subjective experience of pain and activation of SI and ACC is likely attributable to factors other than differential sensitivity of spinal or peripheral afferent mechanisms [which are often show the *same* levels of activity in such cases]”. So, we might add a twist to my case of Mild and Severe: we might stipulate that Mild and Severe’s peripheral and spinal nociceptive neurons fire at the same rates in response to the same noxious temperatures, and that there are only differences in S1 and ACC in the cortex. Given that the only differences are in the cortex, in order to explain why Mild and Severe’s pains differ in intensity, neural projectivists would have to say that Mild and Severe’s pain systems represent different levels of *cortical activity* and project them onto the same bodily regions. And if in this hypothetical case human pain representations only represent cortical neural activity, presumably this is also true in the actual case. The “cortical” version of neural projectivism is also supported by the fact that often the best correlations between the phenomenal and the neural are often to be found in the cortex (e. g. smell similarity correlates best with PPC neural similarity).

Now in any version I think neural projectivism is an interesting attempt to come to grips with the kind of empirical problems I have raised. It would explain how sensory consciousness manages to be both externally-directed and internally-dependent. But it faces challenges that seem to me overwhelming. To illustrate, take the case of pain.

To begin with, some background. The *depth problem* or *distance problem* is a well-known problem for naturalistic theories of sensory representation. Consider a cortical representation of size or orientation or some other spatial property. It is not only causally correlated with an external size. It is also correlated with a pattern of firing on the retina, as well as a pattern of firing in the lateral geniculate nucleus. What

makes it the case that it represents one of these elements in the causal chain, as opposed to the others?

Typically, theories of sensory representation are designed to solve the depth problem in favor of the distal or “distant” properties. Some theories appeal to the notion of function (Dretske 1995). Intuitively, the function of the cortical representation is to indicate size, not some intermediary retinal state. Other theories appeal to “asymmetrical dependence” (Fodor 1990). The cortical representation tracks neural activity only *because* it tracks shape. So, on these theories, the sensory representation represents the external shape. Indeed, these theories seem to entail that *all* of our cortical sensory representation represent distal properties or conditions. So, for instance, a pain representation would seem to have the function of indicating *noxious temperatures*, which are biologically important because they can harm the organism. The function of the pain representation is not to indicate some parameter in fact instantiated in the brain. In general, the function of sensory representations is to indicate biologically significant distal conditions, like level of sugar, stable reflectance properties, and so on. Hence, on standard theories of sensory representation, all cortical sensory representations represent external properties, not properties in fact possessed by neural assemblies.

By contrast, the neural projectivist takes a non-uniform view concerning what properties are represented by atomic sensory representations. He holds that some atomic sensory representations represent distal properties, like location, orientation or shape. But he also holds that other atomic sensory representations represent properties that are in fact only instantiated by neural assemblies, like nociceptive firing rate as opposed to temperature. (Of course, as a projectivist, he doesn’t hold that the sensory system represents these properties *as* instantiated by neural assemblies.) When we have experiences, both sorts of atomic representations are combined into complex representations. The result is a kind of binding error, in which properties that are in fact only instantiated in the head are bound with properties that are instantiated external to the head.

Now I can state my first problem for the view: what is the *general* naturalistic theory of sensory representation, which implies the projectivist’s *non-uniform* answer to the depth problem? As I said, standard theories of sensory representation always solve the depth problem in favor of external properties. I cannot think of alternative general theory of sensory representation that in some cases solves the depth prob-

lem in favor of external properties and in other cases solves it in favor of “internal” properties. In response, the projectivist would at least have to sketch a general view that might imply his non-uniform view, in order to make his view believable.

In reply, the neural projectivist might claim that it would be *good* if the pain system keeps track of nociceptive firing rates, because nociceptive firing rates correlate better with threat level. By contrast, the physical intensity (e. g. temperature) of the external stimulus correlates poorly with threat level. Since it would be good if the pain system keeps track of nociceptive firing rates, it probably has a representation *R* that represents nociceptive firing rates. (See Hill 2012, 138; but note that here Hill is not developing neural projectivism, but the kind of response-dependent intentionalism I will discuss later on in the present section.)

The trouble is that this does not answer my specific challenge. My challenge was: what is the general theory of representation (sensory representation *X* represents *Y* iff . . .) which entails that *R* represents nociceptive firing rates, and at the same time entails that other atomic sensory representations represent external properties like location and size? At best, the reply only provides a *reason to believe* that the pain system has a representation *R* that represents nociceptive firing rates, in addition to other representations that represent external conditions. It does not explain *what makes this the case*, by showing how it follows from general theory of representation.

(It is also worth mentioning that the reasoning in the reply is somewhat questionable. An analogy with shapes may help us to see this. It is of course good to keep track of neural computations of size in the sense that it good to engage in behavior in step with those computations; but this does not provide a reason to think that we have sensory representations, or experiences, that *represent* the relevant computations, as opposed to apparent shapes.)

Here is a second problem for neural projectivism. Even if the neural projectivist can come up with a theory of sensory representation on which some of atomic sensory representations represent internal properties while others represent external properties (thus answering my first problem), he needs a theory of how complex representations manage to represent them as *co-instantiated*. For instance, on his view, when you have a pain in your foot, there is in your brain a complex representation that represents the complex condition *there is something with the properties throbbing, being located in the foot, and exhibiting firing*

rate N . How is this? On many views, any sensory representation represents a condition by co-varying with it under optimal conditions. But this simple view will not work in this case, because the relevant conditions never obtains and doesn't even obtain in nearby counterfactual situations. Of course, the neural projectivist might claim that the content of this complex representation is determined compositionally (like sentences in a language of thought), thanks to a kind of concatenation relation between atomic neural representations. But just when are two atomic neural representations "concatenated". Is there a functional relation of this relation?

Finally, let me mention a third problem for neural projectivism. Given that there is "good internal correlation" and "bad external correlation" in all sense-modalities, if there is a true general theory of sensory representation that entails neural projectivism in the case of pain (contrary to first worry), then it presumably also entails neural projectivism about other experiences. So neural projectivism about pain stands or falls with neural projectivism about other experiences. But, if the view is odd in the case of pain, it is even odder in the case of other experiences. For instance, the idea would be that loudness level is a matter of firing rates of neurons somewhere in the auditory channel (presumably the cortex), but the auditory system projects internal firing rates onto external events. So the content of an auditory experience might be *there is an event has the property of being at place p and the property of exhibiting firing rate N* . Likewise, brightness and hue are neural properties but the visual system projects them into external objects. So the content of a visual experience might be *there is an object that has the property of being at place p , the property of being round, and the property of undergoing neural activity N* . The problems that I raised above against neural projectivism about pain apply with even more force against neural projectivism about other experiences.

Response-dependent intentionalism. The kind of tracking intentionalism that I have taken as my stalking horse holds that the sensible qualities represented by experience are objective properties of external things. By contrast, *response-dependent intentionalism* holds that the sensible qualities represented by experience are *response-dependent properties*. These response-dependent properties might have the form: *causing, or being disposed to cause, internal neural state N in individual or population I* . The idea is that phenomenal character of experience is partly constituted by the representation of such response-dependent proper-

ties. However, the response-dependent intentionalist I am interested in retains the central idea of tracking intentionalism that we represent properties by tracking them. On the basis of actual cases of variation (discussed in §2.1), Uriah Kriegel (2009) has developed this type of view in great detail.

In my view, the best argument for the response-dependent view is that it answers the internal-dependence argument and the structure argument. So, for instance, return to the case of Soft and Loud. On hearing the same tone, they track the same (disjunctive) response-independent physical property of the tone. But on response-dependent intentionalism Soft's auditory experience also tracks and thereby represents the response-dependent property *normally causing firing rate f in the auditory channel of Soft's population*, while Loud's experience tracks and thereby represents the *different* response-dependent property *normally causing firing rate f_+ in the auditory channel of Loud's population*. Hence the neural difference between Soft and Loud is associated with a representational difference. This might explain the phenomenal difference between their experiences in intensity. The same kind of account might be applied to the cases of Yuck and Yum, Mild and Severe, and Sniff and Snort. This would answer the internal-dependence argument.

Response-dependent intentionalism might also avoid my structure argument. So, for instance, on this view, smell qualities might be dispositions to produce ensemble activation states in the PPC. Perhaps these dispositions resemble insofar as their neural manifestations resemble. Then, given good internal correlation, our judgments about resemblances among smell qualities will come out true. Indeed, on this view, if two smell qualities resemble to a degree, they do so essentially, because they are essentially dispositions to produce neural states that resemble do that degree. So this view accommodates the intuition (mentioned in §4.2) that the structural features of sensible qualities are essential to them. Likewise for our judgments about the structural features of other ranges of sensible qualities.

In a way, response-dependent intentionalism resembles neural projectivism. Both views hold that the sensible qualities constitutively involve neural responses. The difference is that, while the neural projectivist holds that the sensible qualities are properties of neural responses erroneously projected onto external items, the response-dependent intentionalist holds that they are dispositions to produce neural responses. Since external items really have those dispositions, the response-dependent intentionalist avoids projective error.

My reply. The main problem with response-dependent intentionalism is that there is no good naturalistic theory of how we represent response-dependent properties in experience. Call this the *psychosemantic problem*.

Kriegel (2009) favors Dretske's (1995) theory of sensory representation. But in fact Dretske's theory of sensory representation is incompatible with Kriegel's response-dependent intentionalism. On Dretske's 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 of the form *normally causing brain state B in humans*. (Even more obviously, it does have the function of indicating Kriegel's more complex response-dependent properties: for reasons I will not go into, he claims that they involve dispositions to produce brain states in all actual creatures, including alien creatures, if such there be.) A brain state has the function of indicating a biologically important property: chemical property, bodily disturbance, etc. So on Dretske's theory of representation, we sensorily represent such response-independent properties, not Kriegel's response-dependent properties. That is of course Dretske's view. Many other theories, for instance Tye's (2000) tracking theory, explain representational relations in terms of *causal* or *explanatory* relations. But the dispositional, response-dependent property normally causing brain state *B* isn't causally efficacious in the production of brain state *B*. So such theories, too, are incompatible with response-dependent intentionalism. In general, I see no theory of sensory representation compatible with response-dependent intentionalism.

Kriegel (forthcoming) answers one objection to response-dependent intentionalism, a kind of *circularity objection* he attributes to Robert van Gulick and Joseph Levine. The circularity problem is simply the problem of characterizing the relevant response-dependent properties in non-phenomenal terms, so that they can be appealed to in an intentionalist theory of phenomenal character without circularity. I agree with Kriegel that this problem can be answered: the responses can be characterized in neural terms, for instance. But my psychosemantic problem remains. Even if external items possess extremely complex response-dependent properties that can be characterized in non-phenomenal terms, the response-dependent intentionalist still needs an account of what makes it the case that our

experiences represent any of them (the standard accounts do not work).¹⁶

Millikan to the rescue? William Lycan (2006) has suggested that the internal-dependence argument might only work against versions of externalist intentionalism that accept a simple tracking theory of sensory representation. Without developing the details, he suggests that Millikan's (1989) more complex consumer-based theory of representation might enable the externalist intentionalist to handle the cases I have discussed.

My reply. I think we can see that this response will not save externalist intentionalism even without going into the details of Millikan's sophisticated consumer-based theory of representation. To illustrate, consider Soft and Loud. They have experiences of the same tone that differ in sensory intensity. What, according to the proponent of a consumer-based theory of sensory representation, is represented by their experiences? As Lycan likes to put it, what are the *representata*? There is no good option consistent with externalist intentionalism.

(i) Maybe Millikan's theory implies that Soft and Loud's experiences represent the response-independent physical properties (involving amplitude, frequency, and critical bands) that constitute the loudness, pitch and timbre of the tone, according to the externalist. So, their experiences have the same content. But this option is inconsistent even with the intentionalist thesis that (at least within a sense modality) experiences differ in sensory phenomenology only if they differ in representational content. The externalist intentionalist needs to find *different representata*. (ii) Another option is that Soft and Loud's experiences represent different firing rates in their own auditory systems, and somehow project these onto the same external sound-event. We have seen the problems with this kind of "neural projectivism". (iii) A final option is that their experiences represent different extremely complex

¹⁶ Suppose Soft and Loud hear the same tone. The tone has a huge set of co-extensive dispositions to cause various neural responses in them and other creatures under various conditions. Even if the response-dependent intentionalist manages to specify a permissive theory of representation on which Soft and Loud sensorily represent such response-dependent properties, he would face a follow-up problem. He does not want to be so permissive as to say that they represent the *same* huge swarm of response-dependent properties, because this would leave him without an account of the phenomenal difference between their experiences. But what could make it the case that Soft sensorily represents one specific response-dependent property within this set and Loud represents a different response-response-dependent property within the set? This might be called the *selection problem* or the *promiscuity problem* (Pautz 2010).

Kriegel-style response-dependent properties. But no theory of representation (including Millikan's) is compatible with this option, as we have seen.

Quasi-intentionalism. Externalist intentionalists are out of options. But maybe they can make a simple retreat. As I defined externalist intentionalism, they are committed to the intentionalist thesis that all phenomenal differences among sensory experiences are constituted by representational differences. Maybe they could reject intentionalism (so understood) and retreat to what I will call *quasi-intentionalism*. On this view, some phenomenal differences among sensory experiences are constituted by representational differences, while others are constituted by merely *functional* or *neural* differences.

Some philosophers have already defended quasi-intentionalism concerning a certain limited range of cases. Thus Lycan (forthcoming) claims that *subtle affective differences* among experiences are not representational differences; they are mere functional differences concerning effects on desire and behavior. So he would reject the representational account of affective phenomenology defended by Cutter and Tye. Likewise Hill (2012, note 2) has said that "how it seems to one to have an experience [e. g. the apparent simplicity of colors] is determined by two factors – it is determined in part by the representational content of the relevant representation, and in part by the representation's intrinsic [neural] and functional properties". If by 'how it seems' he means *phenomenal character* (and not just our inclinations to form sophisticated *beliefs*, e. g. about the simplicity of colors), then Hill endorses quasi-intentionalism about some cases. Some aspects of the phenomenology of experience cannot be explained in terms of representational content, and there are possible cases in which such aspects (apparent phenomenal simplicity perhaps) vary while representational content is held constant. (This *may* not be the correct interpretation because Hill (2009, 148) also says that phenomenal character is nothing but the "set of" *represented properties*, suggesting a one-factor view.)

Maybe quasi-intentionalism could be used to explain the more radical forms of phenomenal variation found in my coincidental variations. The idea is that, in these cases, the individuals involved (Yuck and Yum, Sniff and Snort, Mild and Severe, Soft and Loud) have experiences that differ radically in their sensory character, but they *have exactly the same representational contents*. What constitutes the phenomenal differences are either merely neural differences or merely function-

al differences (e. g. tendencies to group stimuli in certain ways). And maybe the quasi-intentionalist could answer the structure argument by explaining facts about qualitative structure in neural or functional terms, rather than in purely response-independent terms.

My reply. Maybe *some* phenomenal differences are not representational, for instance differences in mood or valence. But I think that the suggested response takes quasi-intentionalism too far, and I think that Lycan would agree. Consider Yuck and Yum, Mild and Severe, Sniff and Snort, and Soft and Loud. Contrary to the suggested quasi-intentionalist response, given that their experiences differ in sensory character, they also *differ* in representational content. The phenomenal differences between them cannot be treated as *merely* neural or functional differences. To see this, let's be fanciful. Suppose you could occupy their points of view on the world, and switch between them. In switching between Soft and Loud the world would *seem* different to you: different loudness-levels would *appear to* attach to the same external sound events. Likewise, the science makes it reasonable to suppose that the experiences of Yuck and Yum, and Mild and Severe, do not just differ in affective valence; they differ in sensory character. So, if you could switch between their points of view, numerically different pain or taste qualities would *appear to be* present in a certain bodily region or in your tongue. So, in some sense of 'representational', their experiences certainly differ in representational content (in what qualities ostensibly present to their subjects), even though they bear the same externally-determined naturalistic relations to the same objective properties.

6. An Edenic Theory of Sensory Consciousness?

Now we have puzzle. Sensory consciousness is both "externally-directed" and "internally-dependent". The individuals in my coincidental variation cases are ostensibly conscious of different qualities "out there", owing to the internal neural differences between them, even though they track the same objective properties. How is this? What in the world are these different qualities? We have seen that they are not objective physical properties (chemical properties, types of damage, and so on). In fact, I would argue that they are not properties of extra-cranial items of any sort. The main alternative to this externalist picture is a traditional internalist picture on which phenomenal types are necessarily identical with internal neural types and the sensi-

ble qualities are neural properties projected onto items in external space. This peculiar view is sometimes called ‘the stinking brain theory’ – or, when color is involved, ‘the colored brain theory’. But we have seen that this kind of projectivist view too faces enormous problems.

In my view, both of these alternatives share a false presupposition, namely that the sensible qualities must be located *somewhere* in the world. The externalists are wrong to locate the sensible qualities outside the brain. But their opponents are also wrong to kick the sensible qualities upstairs into the brain. Although I cannot argue for this here, I would suggest that the overall best view is that, while our brain portrays the world and our bodies as filled with sensible qualities because that enhances adaptive fitness, they are not real qualities that belong to anything, including the brain itself. They are wholly chimerical. This is what David Chalmers (2006) calls the *Edenic theory*.¹⁷ His arguments for the view are based on *a priori* and phenomenological considerations. By contrast, I think that the best argument requires looking at the kind of research in psychophysics and neuroscience I have discussed here.¹⁸

References

- Armstrong, D. (1999). *The Mind-Body Problem*. Boulder: Westview Press.
- Batty, C. (2010). “What’s That Smell.” *Southern Journal of Philosophy*. 321-348.
- Block, N. (1999). “Sexism, Ageism, Racism and the Nature of Consciousness.” *Philosophical Topics* 26: 39-70.
- Block, N. (2010). “Attention and Mental Paint.” *Philosophical Issues* 20: 23-63.
- Block, N. and K. O’Regan. (2012). “Discussion of J. Kevin O’Regan’s *Why Red Doesn’t Sound like a Bell: Understanding the Feel of Consciousness*.” *Review of Philosophy and Psychology* 3: 89-108.
- Borg, G., Diamant, H., Strom, L. and Zotterman, Y. (1967). “The Relation Between Neural and Perceptual intensity.” *Journal of Physiology* 192: 13-20.
- Byrne, A. and D. Hilbert. (2003). “Color Realism and Color Science.” *Behavioral and Brain Sciences* 26: 3-21.
- Byrne, A. (forthcoming). “Hmm . . . Hill on the Paradox of Pain.” *Philosophical Studies*.
- Byrne, A. (2003). “Color and Similarity.” *Philosophy and Phenomenological Research* 66: 641-65.
- Campbell, J. (2002). *Reference and Consciousness*. Oxford: Oxford University Press.
- Chalmers, D. (2006). “Perception and the Fall from Eden.” In T. Szabo Gendler and J. Hawthorne (eds.) *Perceptual Experience*. Oxford: Oxford University Press.
- Chen JY, Victor JD, Di Lorenzo PM. (2011). “Temporal Coding of Intensity of NaCl and HCl in the Nucleus of the Solitary Tract of the Rat.” *Journal Neurophysiology* 105: 697-711.
- Churchland, P. (1996). *The Engine of Reason, The Seat of the Soul*. Cambridge: MIT Press.
- Coghill, R., Sang, C., Maisog, J., and Iadarola, M. (1999). “Pain Intensity Processing within the Human Brain.” *Journal of Neurophysiology* 82: 1934-1943.
- Coghill, R., McHaffie, J. and Yen, Y-F. (2003). “Neural Correlates of Interindividual Differences in the Subjective Experience of Pain.” *Proceedings of the National Academy of Science* 100: 8538-8542.
- Cowart, B.J. and Rawson, N.E. (2001). “Olfaction.” In E. Goldstein (ed.) *The Blackwell Handbook of Perception*. Oxford: Blackwell Publishers.

¹⁷ Interestingly, Ruth Millikan also defends a kind of error theory of sensory experience. In comments on an earlier version of this paper, she asserts that relations among qualities are “chimerical” and do not obtain among any external items, appealing to the work in neuroscience that I discuss. Millikan (this volume) makes remarks along similar lines. For questions about her view here and her argument for it, see Pautz (2011).

¹⁸ I presented the material in this paper at Rutgers University, Columbia University, and the Australian National University. I would like to thank the audiences on those occasions for very helpful comments. My thanks also to Chris Hill and Casey O’Callaghan for some extremely helpful correspondence on the issues I have discussed. Finally, I would like to thank Richard Brown for his work in organizing the *Consciousness Online* conference in which an earlier version of this paper appeared.

- Davies, W. (forthcoming). "Colour Similarities at Different Levels."
- Di Lorenzo PM, Chen JY, Victor JD. (2009). "Quality Time: Representation of a Multidimensional Sensory Domain Through Temporal Coding." *Journal of Neuroscience* 29: 9227–38.
- Dretske, F. (1995). *Naturalizing the Mind*. Cambridge: MIT Press.
- Fodor, J. (1990). *A Theory of Content and Other Essays*. Cambridge: MIT Press.
- Gescheider, G. (1997). *Psychophysics: The Fundamentals*. Lawrence Erlbaum Associates
- Haddad R, Khan R, Takahashi YK, Mori K, Harel D (2008). "A Metric for Odorant Comparison". *Nature Methods* 5: 425–429.
- Hill, C. (2009). *Consciousness*. Cambridge: Cambridge University Press.
- Hill, C. (2012). "Locating Qualia: Do They Reside in the Brain or in the Body and the World?". In S. Gozzano and C. Hill (eds.) *New Perspectives on Tye Identity*. Cambridge: Cambridge University Press.
- Hofbauer, R. K., Rainville, P., Duncan, G. H., & Bushnell, M. C. (2001). "Cortical Representation of the Sensory Dimension of Pain." *Journal of Neurophysiology* 86: 402–411.
- Horgan, T., J. Tienson, and G. Graham. (2004). "Phenomenal Intentionality and the Brain in a Vat." In R. Shantz (ed.) *The Externalist Challenge: New Studies on Cognition and Intentionality*. Amsterdam: de Gruyter.
- Horwitz, G. and C. Haas. (2012). "Nonlinear Analysis of Macaque V1 Color Tuning Reveals Cardinal Directions for Cortical Color Processing." *Nature Neuroscience* 15: 913-919.
- Howard JD, Plailly J, Grueschow M, Haynes JD, Gottfried JA. (2009). "Odor Quality Coding and Categorization in Human Posterior Piriform Cortex." *Nature Neuroscience* 12: 932-939.
- Kalderon, M. (2011). "The Multiply Qualitative." *Mind* 120: 239-262.
- Kenshalo DR, Iwata K, Sholas M, Thomas DA. (2000). "Response Properties and Organization of Nociceptive Neurons in Area 1 of Monkey Primary Somatosensory Cortex." *Journal of Neurophysiology* 84: 719–729.
- Klein, C. (2007). "An Imperative Theory of Pain." *Journal of Philosophy* 104: 517–532.
- Klein, C. and M. Martinez. (forthcoming). "Naturalism and Degrees of Pain."
- Kriegel, U. (2009). *Subjective Consciousness*. Oxford: Oxford University Press.
- Kriegel, U. (forthcoming). In Defense of Self-Representationalism: Reply to Critics. *Philosophical Studies*.
- Langers D, van Dijk P, Schoenmaker E, Backes W. (2007). "fMRI activation in Relation to Sound Intensity and Loudness." *NeuroImage* 35: 709–718.
- Linster, C., B. Johnson, E. Yue, a. Morse, Y. Choi, M. Choi, A. Messiha, M. Leon. (2001). "Perceptual Correlates of Neural Representations Evoked by Odorant Enantiomers." *Journal of Neuroscience* 24: 9837-9843.
- Lycan, W. (2001). "The Case for Phenomenal Externalism." In *Philosophical Perspectives* 15: 17-35.
- Lycan, W. (2006). Pautz vs. Byrne & Tye on Externalist Intentionalism, <https://webpace.utexas.edu/arp424/www/vs.pdf>
- Lycan, W. (forthcoming). "Block and the Representation Theory of Sensible Qualities." In A. Pautz and D. Stoljar (eds.) *Themes from Block*. Cambridge: MIT Press.
- Malnic, B., Hirono, J., Sato, T. & Buck, L. B. (1999). Combinatorial Receptor Codes for Odors." *Cell* 96: 713–723.
- Margot, C. (2009). "A Noseful of Objects." *Nature Neuroscience* 12: 813-814.
- Martinez, M. (2011). "Imperative Content and the Painfulness of Pain." *Phenomenology and the Cognitive Sciences*. 10: 67–90.
- McLaughlin, B. (2003). "Color, Consciousness, and Color Consciousness." In Q. Smith (ed.) *New Essays on Consciousness*. Oxford: Oxford University Press.
- Millikan, R. (1989). "Biosemantics." *Journal of Philosophy* 281-297.
- Moore, B. (2003). *An Introduction to the Psychology of Hearing*. San Diego, CA: Academic Press.
- Noë, A. (2004). *Action in Perception*. Cambridge: MIT Press.
- Oakley B. (1985). "Taste Responses of Human Chorda Tympani Nerve." *Chemical Senses* 10: 469–481.
- O'Callaghan, C. (2002). *Sounds*. Dissertation, Princeton University.
- O'Callaghan, C. (2009). "Auditory Perception." *The Stanford Encyclopedia of Philosophy*. <<http://plato.stanford.edu/archives/sum2009/entries/perception-auditory/>>.
- O'Regan, K. (2011). *Why Red Doesn't Sound like a Bell: Understanding the Feel of Consciousness*. Oxford: Oxford University Press.
- Papineau, D. (2012). "My Philosophical Views." At <http://philpapers.org/profile/10091/myview.html>.
- Pautz, A. (2006a). "Sensory Awareness is not a Wide Physical Relation." *Noûs* 40: 205-240.
- Pautz, A. (2006b). "Can the Physicalist Explain Colour Structure in terms of Colour Experience?" *Australasian Journal of Philosophy* 83: 535-564.
- Pautz, A. (2010). "Do Theories of Consciousness Rest on a Mistake?" *Philosophical Issues* 20: 333-367.
- Pautz, A. (2011). "Reply to Ruth Millikan." <https://webpace.utexas.edu/arp424/www/mh.pdf>
- Price, D. (1999). *Psychological Mechanisms of Pain and Analgesia*. Seattle: IASP Press.
- Price, D. (2002). "Central Neural Mechanisms that Interrelate Sensory and Affective Dimensions of Pain." *Molecular Interventions* 2: 392–403.
- Price DD, Bush FM, Long S & Harkins SW (1994). "A Comparison of Pain Measurement Characteristics of Mechanical Visual Analogue and Simple Numerical Rating Scales." *Pain* 56: 217–226.
- Rainville, Pierre, Duncan, G. H., Price, Donald D., Carrier, M., and Bushnell, M. Catherine. (1997). "Pain Affect Encoded in Anterior Cingulate but Not Somatosensory Cortex." *Science* 277: 968–71.
- Relkin EM and Doucet JR. (1997). "Is Loudness simply Proportional to the Auditory Nerve Spike Count?" *Journal of Acoustical Society of America* 101: 2735–2740.
- Röhl M, Kollmeier B, Uppenkamp S (2011). "Spectral Loudness Summation Takes Place in the Primary Auditory Cortex". *Human Brain Mapping* 32: 1483–1496.
- Sakurai, T., Misaka, Y., Ueno, M., Ishiguro, S., Matsuo, Y., Ishimaru. (2010). "The Human Bitter Taste Receptor, hTAS2R16, Discriminates Slight Differences in the Configuration of Disaccharides." *Biochem Biophys Res Commun* 402: 595–601
- Schiffman, S.S. and Erickson, R.P. (1971). A Psychological Model for Gustatory Quality. *Physiology and Behavior* 7: 617-633.
- Shoemaker, S. (1994). "Phenomenal Character." *Noûs* 28: 21-38.
- Shoemaker, S. (2000). "Phenomenal Character Revisited." *Philosophy and Phenomenological Research* 60: 465-467.
- 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.
- Smith, B. (2007). "The Objectivity of Tastes and Tasting." In B. Smith (ed.) *Questions of Taste*. Oxford: Oxford University Press.
- Stevens, S. S., Carton, A. S., & Schickman, G. M. (1958). A Scale of Apparent Intensity. *Journal of Experimental Psychology* 56: 328–334.
- Timmermann L, Ploner M, Haucke K, Schmitz F, Baltissen R, Schnitzler A (2001). "Differential Coding of Pain Intensity in the Human Primary and Secondary Somatosensory Cortex." *Journal of Neurophysiology* 86: 1499–1503.
- Turin, L. (2002). "A Method for the Calculation of Odor Character from Molecular Structure." *Journal of Theoretical Biology* 216: 367–385.
- Tye, M. (2000). *Consciousness, Color and Content*. Cambridge: MIT Press.
- Tye, M. (2006). "The Puzzle of True Blue." *Analysis* 66: 173–178.
- Tye, M. (2009). *Consciousness Revisited*. Cambridge: MIT Press.
- Tye, M. and B. Cutter. (2011). "Tracking Intentionalism and the Painfulness of Pain." *Philosophical Issues* 21: 90-109.
- van der Heijden, A. (1993). "Sweet and Bitter Tastes." In T. E. Acree & R. Teranishi (eds.), *Flavor Science: Sensible Principles and Techniques*. Washington, DC: ACS Books.
- Walters D. (1996). "How are Bitter and Sweet Taste Related?" *Trends in Food Science and Technology* 7: 399–403.
- Youngentob, S.L., Johnson, B.A., Leon, M., Sheehe, P.R. & Kent, P.F. (2006). "Predicting Odorant Quality Perceptions from Multi-dimensional Scaling of Olfactory Bulb Glomerular Activity Patterns." *Behavioral Neuroscience* 120: 1337–1345.

