

DO THEORIES OF CONSCIOUSNESS REST ON A MISTAKE?

Adam Pautz
University of Texas at Austin

I'm always sympathetic to Bishop Butler's 'Everything is what it is and not another thing'

–Saul Kripke

I think standard theories of consciousness rest on a mistake. They assume *reductionism*: states involving phenomenal consciousness are physical-functional states. But I think consciousness resists reductive explanation. This is *primitivism*.

The name is imperfect. It suggests a spooky dualism. It also suggests a no-theory theory. But primitivism is compatible with physicalist supervenience and naturalism about the mind. It is also compatible with interesting *non-reductive* theories of consciousness. I will sketch such a theory here.

Primitivism is not preposterous. Reductionism is not science but speculative metaphysics. The standard (simplicity and causal) arguments for it are weak. Even among naturalistically inclined philosophers, primitivism is taken very seriously for reference, the intentionality of thought, goodness, color, and knowledge. Why not for consciousness?

To explain why I favor primitivism, I must distinguish between two questions:

The general question: What distinguishes states that exhibit phenomenal consciousness from states that do not?

The quality question: Among phenomenally conscious states, what determines differences in quality?

Answers to the first question include first-order accessibility theories, higher-order theories, self-representational theories, and biological theories. Answers

to the second question include biological theories (quality differences are necessarily differences in the neural content-carriers) and intentional theories (quality differences are generally differences in the contents carried).

The focus has been on the general question. For instance, Prinz defends a first-order accessibility theory.¹ But what about the quality question? The external-directedness of consciousness suggests intentionalism. Differences among conscious experiences reside in differences in the sensible properties (colors, pains, tastes, sound-qualities, etc.) represented. So an answer to the quality question must include a reductive theory of all sensible properties and a reductive theory (psychosemantics) of how we represent them. Prinz accepts intentionalism but mostly neglects these issues (but see §5). Yet, on intentionalism, since nearly every conscious experience is necessarily intentional, a reductive answer to the general question requires a reductive theory of sensory intentionality and all sensible properties.

My case for primitivism involves the quality question. I accept intentionalism but think there is no good *reductive* theory of the sensible properties or of how we represent them. Those who focus on the general question and ignore the quality question have missed the most difficult part of the puzzle of consciousness.

My arguments for primitivism are unique. First, I rely on empirical considerations, not Zombies, Revelation, or Mary. Second, I focus on non-visual modalities. My main target is the combination of intentionalism and a broadly ‘tracking’ psychosemantics of the kind defended by Fodor, Dretske, Neander, Stalnaker, Tye and others (§§1–4).² Along the way I introduce constraints on a theory of consciousness. Then I develop an overlooked problem with Shoemaker’s theory (§5). What emerges is a systematic argument that no version of reductionism can simultaneously satisfy all the relevant constraints. Finally, I show that primitivism can do so without difficulty (§6).

1. Consciousness as Externally-Directed: The Case for Intentionalism

Since my arguments presuppose intentionalism, I begin by arguing for it. I reject the transparency argument. Instead, my argument depends on the external-directedness of consciousness.³

Suppose you see a tomato. There is a salient property you have in this case and all possible cases in which you have a phenomenally identical experience. Call it *r*. And call such properties *conscious properties*, since they involve phenomenal consciousness.

Intuitively, *r* is not like the mere sequence of marks ‘a red and round object is present’, which has its externally-directed properties contingently. Instead:

External-Directedness: R has at least *some* (second-order) externally-directed properties *necessarily*.

For instance, R necessarily possesses the *externality property*: being such that, if one has it, then one has an experience as of a red_p and round object at a certain distance and viewer-relative location. Here ‘round’, ‘distance’, ‘location’ have their usual meanings. I use ‘red_p’ for sensible redness and wish to remain neutral on whether it is an ordinary color or a Shoemakerian appearance property (§5). Further, R necessarily possesses the *matching property*: being such that, if one has it, then one is in a state that matches the world only if a red_p and round object is present. It necessarily possesses the *grounding property*: being such that, if a believer has it for a non-trivial period, then he thereby has the capacity to have a belief that matches the world only if such an object is present. Finally, it necessarily possesses the *justification property*: being such that, if a believer has it in the absence of defeaters, then he has a justification for having such a belief. Even if R might occur in a brain in a vat or be normally caused by green squares (which many externalists deny), it cannot be stripped of these properties.

The biological theory of Block and McLaughlin holds that R is necessarily identical with some internal (perhaps global) neural property N .⁴ External-directedness undermines this theory:

- 1 R has *some* externally-directed properties necessarily: for instance, it is necessarily as of a round thing at a certain distance and viewer-relative location.
- 2 By contrast, N (like the sentence ‘a red and round object is present’) has *all* its externally-directed properties only *contingently*, owing to connections with external properties and action.
- 3 Hence, by Leibniz’s Law, R cannot be necessarily identical with N .

The case for 2 is that in some possible cases N is stripped of its relations to external properties and action, and hence is stripped of its externally-directed properties (e. g. a world in which N occurs in a lone Petri dish). Since R by contrast cannot be stripped of these properties, R is not necessarily identical with N . Briefly, since experience, like thought, is necessarily externally-directed, a purely biological theory of experience is no more plausible than a purely biological theory of thought.

The argument generalizes. Let P , T , S be pain, taste, and sound conscious properties, respectively. Maybe these conscious properties, unlike R , do not necessarily present an *object*. But P and T necessarily possess analogous externally-directed properties with respect to certain *qualities* and a *place of the body* (in the case of T , the tongue). And S necessarily possesses these

properties with respect to certain events, qualities like loudness and pitch, and temporal and spatial relations and properties. Since the connection between the conscious properties and a certain ostensible bodily or external spatiotemporal region is necessary, but the connection between any neural property N and such a region is only contingent (deriving from N 's contingent relations to the region), these conscious properties cannot be necessarily identical with any neural property N .

Block accepts *biological intentionalism* (personal communication). McLaughlin considers it a possibility.⁵ On one version, a conscious property c is necessarily identical with a neural property N , and necessarily results in representing a content involving external (e. g. spatiotemporal) properties. This entails that N , like c , necessarily results in representing a content involving external properties, and hence has certain externally-directed properties necessarily, contrary to premise 2 above.

But this yields a *magical view of intentionality*: that in every possible situation (hence *even in situations in which it is stripped of its actual links to action and external properties*) N (that is, on this view, c) magically represents a content involving certain external (e. g. spatiotemporal) properties. This is like saying the marks 'a red and round object is at place p ' necessarily represent a red, round object at p *regardless of their use in language*. No externalist or internalist⁶ reductive account of the representation relation (what I call the conscious-of relation below) supports this view. So this magical view is unavailable to biological theorists, who want a reductive account of the mind. Block (personal communication) suggested that Shoemaker has provided an account of sensory representation supporting this view, thereby making it non-magical. But this is incorrect (§5). In any case, as a functionalist, Shoemaker denies the claim of biological intentionalism that there is a neural property (even a 'total realizer') N that is both necessarily identical with c and that has its externally-directed properties necessarily.

Disjunctive theories, sensorimotor theories, and sense datum theories (at least those on which sense data are three-dimensional) accommodate External-Directedness but fail for other reasons. The best option is to accept an ordinary intentional theory. On a simple *property-complex version*, conscious properties are relational properties of the form: standing in intentional relation R to a certain cluster of external (possibly uninstantiated) properties P , for instance *being bitter and at the tip of my tongue*.⁷ I will call this relation the *conscious-of relation* or the *sensory representation relation* and the properties in its range *sensible properties*. Since such "big", world-involving relational properties, unlike neural properties, have (as it might be, tracking) relations to external (e. g. spatiotemporal) properties *built in*, they have certain externally-directed properties necessarily. In having built-in external-directedness, they are like conscious properties. Hence, unlike neural properties, they are fit for identification with conscious properties.

2. The Internal-Dependence Argument Against Tracking Intentionalism

Tracking intentionalism is the most natural reductive form of intentionalism. It has two parts. The first is *response-independent reductionism* about the sensible properties. Pains are bodily disturbances, tastes are chemical properties, and sound-qualities are physical properties of sound-events. The second part is a *tracking theory* of how we become conscious of these external properties. In Tye's version, the conscious-of relation is:

The optimal cause relation: x is in a state that has the consciousness-making property and that would be caused by (track) the instantiation of external property y were optimal conditions to obtain.⁸

The "consciousness-making property" is a place-holder for the tracking theorist's answer to the general question: what turns mere representation into conscious presentation. It might be accessibility, or being the actual or potential target of a higher-order thought, or being self-represented. Thus, the conscious-of relation might involve an internally-directed component in addition to an externally-directed component. The tracking intentionalist answers the quality question by adverting to the external property tracked under optimal conditions. For evolved creatures, these are conditions in which our sensory systems are operating in accordance with design.

Likewise, Dretske has an *indicator psychosemantics*. I will assume, as Dretske does, that this theory also goes with response-independent reductionism about the sensible properties (but see §5). Then it yields a version of tracking intentionalism subject to my arguments.⁹

I have three arguments against tracking intentionalism. The arguments of the present section and the next depend on empirical facts. Tracking intentionalism entails that phenomenology is fully determined by the physical property optimally tracked. The whole history of psychophysics and neuroscience goes against this.

Two empirical facts are relevant. First, psychophysics reveals that, even under optimal conditions, there is an extremely bad correlation between phenomenology and the physical property tracked. Second, neuroscience reveals that phenomenology is much better correlated with neural firing patterns in the brain.

For instance, on tracking intentionalism, pain intensity is presumably fully constituted by the intensity and area of the bodily disturbance optimally tracked. But pain intensity is very poorly correlated with these factors. In particular, 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. For instance, doubling electric shock far more than doubles pain intensity. Similarly for heat-induced




Figure 1. Individual psychophysical pain intensity curve in subjects injected with subcutaneous ascorbic acid (A) and normalized fMRI signal (B). From Porro et al. (1998).

pain. Furthermore, pain intensity is a product of stimulus area as well as stimulus intensity, *in a way that cannot be codified*.¹⁰

The explanation is that, even under optimal conditions, there is an uncodifiable power function relationship between cortical somatosensory firing rates and the external bodily disturbance. And studies show that these firing rates are in turn “linearly related to subjects’ perceived pain intensity” (Coghill et al. 1999, 1936). In other words, while there is bad external correlation, there is very good internal correlation (Figure 1).

Likewise for taste and smell. There is bad external correlation. Two sweet tastes might be very different, but optimally track molecules with very similar structures. Conversely, two bitter tastes might be very similar, but optimally track molecules with very different structures. Likewise for smell (see Figure 2).

But good internal correlation obtains. Resemblances among taste experiences are remarkably well-correlated with resemblances among *across-fiber patterns*—patterns of activation across taste-related neurons (see Figure 3). Further, taste intensity is linearly related to neuronal firing rates in the chorda tympani. Similarly for smell.¹¹

Finally, in addition, the relationship between perceived intensity and the amplitude optimally tracked is one of *uncodifiable response compression*. Doubling perceived intensity requires far more than doubling amplitude. At




Figure 2. The smell experiences optimally caused by the chemical properties in row (b) resemble each other more than those optimally caused by the chemical properties in row (a). But it is not the case that the chemical properties in row (b) resemble each other more than those in row (a). From Goldstein (2007).

lower amplitudes, loudness increases more rapidly with increasing amplitude. In addition, 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 critical bandwidths. Likewise the relationship between perceived pitch and the frequency (for complex sounds, fundamental frequency) optimally tracked is one of response compression, with the degree of compression depending on frequency level. Perceived pitch also depends on amplitude. 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.¹²

While there is bad external correlation, there is some evidence of good internal correlation. Perceived pitch appears to depend on some combination of place coding and temporal coding. And loudness is presumably well correlated with the total neural activity evoked by a sound-wave. Loudness models incorporating this assumption have been successful, although it has not been directly confirmed.¹³

Bad external correlation and good internal correlation suggest, against tracking intentionalism, that phenomenology is not fully determined by the external physical property optimally tracked. It is also shaped by




Figure 3. Axes represent activity in taste-related neurons. Across-fiber pattern similarities (here represented by distance) correlate well with taste similarities. From Smith et al. (1983).

postreceptor neural processing. My first argument against tracking intentionalism centers around counterexamples that make the conflict vivid.

The best counterexamples, in my view, concern *hypothetical coincidental variation cases*.¹⁴ In such cases, two individuals optimally track *exactly the same* external properties (F, G, \dots), but they do so *via* different neural states (P, Q, X, Y, \dots), which stand in different similarity-difference relations, and lead to systematically different sorting and other behaviors (see Figure 4).




Figure 4. Coincidental variation.

First case. Mild and Severe occupy a nearby possible world and belong to different species occupying different environments. Each is subject to a lesion in his arm. In Mild, the lesion is harmless. So under optimal conditions it produces a mild somatosensory firing rate. By contrast, the species Severe belongs to evolved in an environment in which such a lesion is extremely dangerous, because in this environment a certain sort of airborne bacterium is present in abundance. Therefore, under optimal conditions, the lesion produces a somatosensory firing rate twice greater than that produced in Mild. As a result, Severe attends to his leg with greater urgency than does Mild. Further, Mild and Severe have different internal similarity metrics for the same external bodily disturbances. Two bodily disturbances, which optimally produce very different somatosensory states in Severe, might optimally produce very similar somatosensory states in Mild, and *vice versa*. In consequence, they treat the same disturbances very differently.

This is the case described in non-phenomenal terms. The question is: is the right verdict Different Experiences or Same Experiences?

Two arguments suggest Different Experiences. First, in humans and other animals there is linear correlation between pain intensity and somatosensory firing rates, whereas there is no codifiable relationship between pain intensity and the bodily disturbance optimally tracked. So the differences between Mild and Severe in somatosensory firing rates are *very good evidence that they have pains of different intensities*, and the fact that Mild and Severe optimally track the same bodily disturbance is *very poor evidence that they have the same pain experience*. Second, the action-oriented differences between Mild and Severe support the same verdict. This is so on functionalist and sensorimotor theories of consciousness, but the point is independent of any theory.

However, despite the vast neural and behavioral differences between Mild and Severe, tracking intentionalism implausibly entails Same Experience, for by stipulation their somatosensory states, although different, have the function of indicating, and are optimally caused by, exactly the same type of bodily disturbance.

Second case. Yuck and Yum occupy a nearby possible world and belong to different species occupying different environments. Some berries

are extremely poisonous to Yuck. But, in Yum's environment, the berries are a very important foodsource, since other foodsources are scarce. So Yum's species evolved immunity to the berries. Further, when Yuck and Yum taste the berries, their taste systems undergo radically different across-fiber patterns. They are also innately disposed respond to the berries with radically different behaviors. For instance, Yuck vomits and withdraws from the berries violently, while Yum is drawn to them, rubs his tummy, and so on. Yet we can stipulate that Yuck and Yum are the same at the receptor-level, so that, when they taste the berries, their postreceptoral across-fiber patterns, although different, optimally track the very same chemical property c of the berries. This case resembles actual cases.¹⁵

In addition, suppose that the poison dart frog is highly poisonous to both Yuck and Yum. When Yuck tastes berries, the across-fiber pattern he undergoes is quite similar to the one he undergoes when he tastes the frog. By contrast, when Yum tastes the berries, the across-fiber pattern he undergoes is radically different from the one he undergoes when he tastes the frog, and much more like the one he undergoes when he tastes nutritious bananas.

Again, Different Experiences is the reasonable verdict. In humans and other animals, taste quality is much better correlated with across-fiber patterns than with the chemical properties optimally tracked. Therefore, the across-fiber pattern differences between Yuck and Yum are very good evidence of phenomenal differences, while the sameness in the chemical properties optimally tracked is very poor evidence of phenomenal sameness. More specifically, in view of the differences in the internal structures of their across-fiber patterns, Yuck probably has *phenomenally similar* taste experiences of the poisonous dart frog and the berries. Whereas Yum has *phenomenally different* taste experiences of these, and his taste experience of the berries is phenomenally similar to taste experience of bananas.

But tracking intentionalism implausibly entails that Yuck and Yum have exactly the same taste experiences of these foodstuffs, despite their vast behavioral and neural differences, because their across-fiber patterns, although different, optimally track the very same external chemical properties.

Third case. Loud and Soft belong to different species. Their auditory systems are alike at the receptor-level, and so optimally track exactly the same physical properties (involving frequency and amplitude) of some sound-event. But, since the sound-event is a mating call to Loud but not to Soft, these same physical properties optimally produce radically different postreceptoral processing in Loud and Soft related to loudness and pitch. The result is that, by every psychophysical test (involving discrimination, grouping, etc.), they have different sound experiences of this sound event and others. Again, they probably have different sound experiences, but tracking intentionalism delivers the opposite verdict.

Since one counterexample would suffice, undermining tracking intentionalism only requires:

Internal-Dependence: In at least *one* possible coincidental variation case, Different Experiences obtains.

The official argument is:

- 1 By tracking intentionalism (Tye, Dretske, Hill), in *every* possible coincidental variation case, Same Experience holds.
- 2 But Internal-Dependence is more reasonable.
- 3 So tracking intentionalism fails.¹⁶

Two clarifications. (i) Internal-Dependence is *not* Internalism: internal, neural factors *completely* determine experience. It only says: certain internal differences, when accompanied by suitable functional-sensorimotor differences, are accompanied by phenomenal differences. Internal-Dependence is consistent with non-neural factors in the first instance determining experience. Maybe neural factors are relevant to experience because they are relevant to world-involving sensorimotor skills or functional role. (ii) Internal-Dependence does not require forsaking intentionalism and External-Directedness and adopting the biological theory (§§5–6). It is just inconsistent with tracking intentionalism's way of accommodating External-Directedness.¹⁷

3. The Structure Argument Against Tracking Intentionalism

My second argument is based on bad external correlation and concerns *structure judgments* such as:

- [1] pain experience P_2 is roughly twice greater than pain experience P_1 .
- [2] Taste experience E_1 is more like E_2 than E_3 .
- [3] Pineapple smell experiences P_1 and P_2 resemble each other more closely than other smell experiences O_1 and O_2 resemble each other.
- [4] The intensity of sound experience S_2 is roughly twice greater than that of S_1 .
- [5] The intensity of sound experience S_1 equals that of S_2 (although they are optimally produced by sound events with different amplitudes and frequencies).
- [6] The pitch-difference between sound experiences S_1 and S_2 is roughly the same as that between S_2 and S_3 .

Introspection and psychophysics support:

Structure: at least *some* such structure judgments are true.

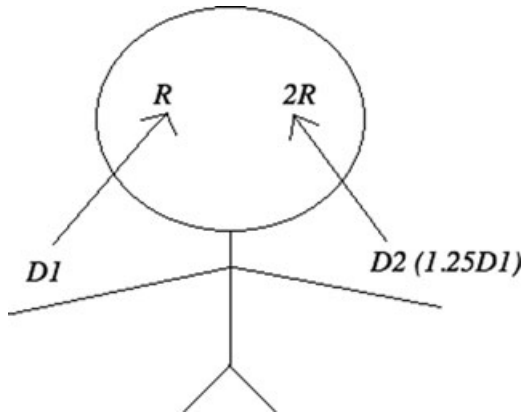


Figure 5. Maxwell.

Now, on tracking intentionalism, [1]–[6] will be true only if the response-independent, physical properties (chemical properties, properties involving amplitude and pitch) *represented by* (optimally tracked by) these experiences stand in the relevant structural relations. And bad external correlation (uncodifiable response expansion or compression) means that they do not (see Figure 2). This yields the Structure Argument:

- 1 If tracking intentionalism is true, then many structure judgments along the lines of [1]–[6] are false.
- 2 But Structure is more reasonable.
- 3 So tracking intentionalism fails.

To support premise 1, I will focus on [1]. Suppose Maxwell undergoes bodily disturbance D_1 in his right arm, which under optimal conditions produces somatosensory firing rate R (see Figure 5). Then he undergoes bodily disturbance D_2 in the corresponding part of his left arm. By any natural measure, D_2 is much less than twice greater than D_1 . Maybe, for instance, it is only 25% greater. But, due to response expansion, the somatosensory firing rate D_2 produces in Maxwell is double that produced by D_1 .

This case is schematic but there are actual such cases involving, for instance, shock and noxious temperature. Since pain intensity correlates with somatosensory firing rates, Maxwell will report [1]. But tracking intentionalism entails that [1] is false:

- 1 On tracking intentionalism, phenomenal facts are in the first instance facts about the properties represented by experience (the properties to which we bear the conscious-of relation), not facts about the neural content-carriers.

- 2 So, on tracking intentionalism, [1] is true if, and only if, the *pain-property* represented by P_2 is roughly twice greater than the *pain-property* represented by P_1 .¹⁸
- 3 Since optimal conditions obtain, on tracking intentionalism, these are the actual disturbances D_2 and D_1 .
- 4 So, on tracking intentionalism, [1] is true iff D_2 is roughly twice greater than D_1 .
- 5 But D_2 is not roughly twice greater than D_1 — D_2 is *much less than* twice greater than D_1 .
- 6 So, on tracking intentionalism, [1] is *false*.

Premise 5 is true even if D_2 and D_1 are firing rates of peripheral nociceptive fibers, as opposed to shocks or noxious temperatures. For we can suppose that response expansion does not take place peripherally but further downstream.

To defuse the Structure Argument, it would not be enough to say [1] is false on the grounds that (contrary to psychophysical research) ratio-scaling of pains makes no sense. Remember: any of the structure judgments [1]–[6] could be used to illustrate the argument. It would not do to say that all are false. A general response is required. I will address three. Each faces a *psychosemantic problem*.

First, the *misrepresentation response*. Hill does not consider my argument, but he makes remarks that suggest a response.¹⁹ On this response, contrary to premise 3, Maxwell's pain experiences P_2 and P_1 do not represent the actual disturbances D_2 and D_1 , where D_2 is much less than twice greater than D_1 . Rather, they represent bodily disturbances that do *not* occur in his arms, such that the second is (in some good sense) twice greater than the first. Hence, his *first-order* representation of his body is inaccurate. However, his *second-order*, introspective judgment [1] about his pain experiences is a true judgment about these represented, non-actual bodily disturbances.

This requires that Maxwell's pain system either *overestimates* the disturbance in his left arm or else *underestimates* the disturbance in his right arm:

- (i) Maxwell accurately represents D_1 in his right arm and inaccurately represents D_2 -plus (D_2 -plus $>$ the actual disturbance D_2) in his left arm, where D_2 -plus (unlike D_2) is twice greater than D_1 .
- (ii) Maxwell inaccurately represents D_1 -minus (where D_1 -minus $<$ the actual disturbance D_1) in his right arm and accurately represents D_2 in left arm, where D_2 is twice greater than D_1 -minus.

But we assumed optimal conditions obtain: there is no top down influence, hyperalgesia or hypoalgesia, temporal or spatial summation, or other interference. Response expansion is the normal, default mode of the pain-system. Hence, in accordance with premise 3, the kind of

psychosemantics favored by Tye, Dretske and Hill is incompatible the misrepresentation entailed by (i) and (ii). But the *main* problem with the misrepresentation response is that there would appear to be no naturalistic facts that could determine *which* of (i) and (ii) obtains. This arbitrariness problem goes beyond the general problem posed by normal misrepresentation for psychosemantic theories.²⁰

Second, the *strange properties response*. It resembles Tye and Bradley's reply to the problem of *color structure*, which attempts to correct for bad external correlation with non-linear functions.²¹ It has two parts.

First part. There is a function, f , from bodily disturbances onto numbers such that, if bodily disturbances D_2 and D_1 optimally cause pain-experiences P_2 and P_1 , and P_2 is judged roughly twice greater than P_1 , then $f(D_2)$ is roughly twice greater than $f(D_1)$. Say that a bodily region has an "ouch value" of N just in case the value of f for the type of bodily disturbance there is N . And call properties of the form *having a value of N* 'strange properties'. Although we ascertained the function f by appealing to our responses, the function itself is a response-independent abstract object. However, it resists response-independent codification, since (even under optimal conditions) the relationship between bodily disturbance and pain intensity differs for different types of bodily disturbance, and is a joint product of stimulus intensity and size.

Second part. Contrary to premise 3, Maxwell's pain experiences P_2 and P_1 do not represent D_2 and D_1 . Rather, they represent the relevant bodily regions as *having an ouch value of 10* (say) and *having an ouch value of 20*. The tracking intentionalist can then say that Maxwell's introspective judgment [1] is a true report about these ouch-values.

The response generalizes. For instance, it generalizes to auditory structure judgments. Again, since (even under optimal conditions) perceived loudness and pitch are related to frequency, amplitude, and critical bandwidths in a messy way, the relevant functions resist codification.

The problem is that this response is not consistent with any reasonable psychosemantics. There are *infinitely many* functions like f and hence infinitely many properties of Max's arm of the form *exhibiting a bodily disturbance D such that $f(D) = N$* . Maxwell's brain state neither is optimally caused by, nor has the function of indicating, any *one* of them to the exclusion of the others. Rather, it bears these naturalistic relations to his arm's being subject to certain non-number-involving worldly bodily disturbances D_2 and D_1 . (Analogy: our brain states track and represent unit-free sizes, not numeric properties like *having a volume in cubic inches of N* .) The problem is general. I can see how, when an ordinary person talks about water, he is talking about H_2O . But I cannot see how, when an ordinary person talks about the loudness-relations among sounds, or resemblances among smells, the truth-conditions of what she says involve amazingly complicated functions which practitioners of psychophysics can (at best) only gesture at.

Third, there is the similar *strange relations* response (Michael Tye in discussion). Again, it has two parts.

First part. There exists a relation R among bodily disturbances with the following property: if bodily disturbances D_2 and D_1 optimally cause pain experiences P_2 and P_1 , and P_2 is judged roughly twice greater than P_1 , then D_2 stands in physical relation R to D_1 . Now, the degree response expansion differs for different types of bodily disturbance. And pain intensity depends on stimulus size as well as stimulus intensity. Thus R is an uncodifiable infinitely disjunctive relation: roughly, D_2 and D_1 stand in R iff D_2 and D_1 are electric shocks applied to the skin with so-and-so spatial extents and D_2 is 25% greater than D_1 , or D_2 and D_1 are noxious temperatures with such-and-such spatial extents and D_2 is 33% greater than D_1 , or . . .

Second part. Contrary to the strange properties response, premise 3 is true. So is premise 4. But the correct interpretation of 4 is not the face-value one. We might have thought that there is a single relation picked out by ‘is twice greater than’ that obtains among worldly determinate length-properties, mass-properties, and so on. Assignments of numbers to properties are valid when they mirror those properties’ objective ratio-structure. But, on the strange relations response, when we use ‘is twice greater than’ in connection with *pains*, as in [1], we are picking out the quite different relation R . This yields non-uniform semantics for ‘is twice greater than’. We might use ‘is twice_p greater than’ for R . So premise 5 is false. Maxwell’s judgment is true iff D_2 is roughly twice_p greater than D_1 . And it is. So, given a non-uniform semantics for ‘is twice greater than’, tracking intentionalism accommodates Maxwell’s structure judgment [1]. Likewise for [2]–[6].

Again, the problem is that there is no credible account of how the truth-conditions of [1]–[6] might involve such strange relations. In reply, the tracking intentionalist might invoke the principle of charity and a Lewisian “use-plus-eligibility” theory to explain how the truth-conditions of [1]–[6] involve such strange relations.²² But such a theory would pull the rug out from under his feet. For it suggests the following empirically-based, semantic argument *against* tracking intentionalism and for the biological theory (§1):

- 1 Phenomenology is the subject of [1]–[6].
- 2 Given bad external correlation and good internal correlation, and reductive physicalism, the most eligible candidates to be the subject of [1]–[6] are *neural properties and their relations*.
- 3 So, the biological theory is probably right.

Call this the *Structure Argument for the biological theory*. The case for 2 is straightforward. On the biological theory, by contrast to tracking intentionalism, there is no need to appeal to wildly unnatural strange relations. In [1], ‘twice greater than’ has its customary meaning, and is being accurately applied to the neuronal firing rates which constitute the intensities

of Maxwell's pains (see Figure 2). Likewise [2]–[6] simply accurately report structural relationships among neural states.²³

I think the Structure Argument for the biological theory must be unsound because that theory violates External-Directedness (§1). But I wish to point out in passing that it may be the best argument for the biological theory, although biological theorists do not use it. It would help the biological theorist answer Block and Papineau's harder problem of consciousness.²⁴ One aspect of the problem is that in humans every conscious property is presumably coextensive with both a "narrow" neural property and a "wide" property of the form *bearing the optimal cause relation to external property P*. So the coextension facts among humans are compatible with both the biological theory and tracking intentionalism. How then might we decide whether the conscious property is identical with the narrow property or the wide alternative? The standard Hill-McLaughlin argument for the biological theory seems to fall short, since it only relies on facts about coextension.²⁵ By contrast, the Structure Argument also relies good internal correlation (match of structural relations) and bad external correlation, together with semantic considerations. These considerations tip the scales in favor of the "narrow" candidate. In short, the biological theory easily accommodates Structure. It also accommodates Internal-Dependence.

But, again, the biological theory violates External-Directedness. An adequate theory must accommodate External-Directedness as well. Before turning to such theories, a final argument against tracking intentionalism.

4. The Percipi Argument Against Tracking Intentionalism

Suppose we are discussing a severe pain in my leg. Intuitively:

[7] Necessarily, if *it* occurs, then someone has a pain experience (focal or nonfocal).

Tracking intentionalism entails that this *Percipi intuition* [7] is false. For, given the context, 'it' cannot refer to the pain experience, which on tracking intentionalism is presumably in the brain. My attention is focused out, on the represented property. On tracking intentionalism, this is nothing but a type of bodily disturbance *D*. On tracking intentionalism, then, [7] is false, because *D* could occur without *anyone having any kind of pain experience at all*, for instance in an insentient cadaver. On this view, 'it' is like a color, which can 'occur' (something can have it), without anyone having an experience.

In reply to a similar point by Block, Tye has said

in a world without experiencers, there can be no pain but there certainly can be tissue damage. [The reason is that] we apply the term 'pain' to tissue damage only

in a certain context—the context provided by tissue damage being represented by a [certain pain experience] (2006, 166).

For Tye, ‘pain’, when used to describe the property represented, is short for a description, ‘the property represented by pain experience *P*’. This indeed validates some of our intuitions, for instance:

[8] Necessarily, if there is (severe) pain in a bodily region, there is a pain experience.

For, on Tye’s descriptivist view, this is equivalent to:

[8*] Necessarily, if there is the property represented by pain experience *P* in a bodily region, then there is a pain experience.

This is true under a narrow scope, *de dicto* reading.

But I chose my original example-sentence [7] to avoid Tye’s response to Block. For [7] does not contain ‘pain’. So even if some occurrences of ‘pain’ are abbreviated descriptions, the descriptivist, *de dicto* maneuver does not apply to [7]. [7] is unambiguously *de re*, and, on tracking intentionalism, could only be about *D*. So, on tracking intentionalism, [7] is false. This is unacceptable.

5. Shoemaker’s Response-Dependent Intentionalism: An Overlooked Problem

The biological theory accommodates Internal-Dependence and Structure but not External-Directedness (§1). Tracking intentionalism (§§2–4) faces the reverse problem: it accommodates External-Directedness but not Internal-Dependence or Structure (or Percipi). A major part of the puzzle of consciousness is finding a theory that simultaneously accommodates all these constraints.

Shoemaker’s motivations are more *a priori* than mine.²⁶ But he ends up with the kind of intentionalism we seek. It goes as follows. Experience involves both qualia and content. For humans, qualia are at least partly, perhaps completely, definable in functional terms. The content of an experience attributes to an object or bodily region a *response-dependent* property of roughly the form: causing, or normally causing, so-and-so experience. This contrasts with tracking intentionalism’s *response-independent reductionism* about the sensible properties. Finally, according to Shoemaker, necessarily, if one has an experience with a certain quale (that is to say, on Shoemaker’s view, if one is in a certain functional state), then one has an experience that represents a certain response-dependent property.

Shoemaker's intentionalism accommodates External-Directedness. It also accommodates Internal-Dependence. Since Shoemaker holds that sensory representation is somehow internally-determined, and the individuals in coincidental variation cases are in different internal-functional states, he would say that they bear the sensory representation relation to different response-dependent properties of the form *normally causing internal state S*. Shoemaker's intentionalism also accommodates Structure, even if in accordance with intentionalism the structure judgments [1]–[6] are about the properties of external objects or bodily regions represented by our experiences. For, on Shoemaker's intentionalism, these are *response-dependent* properties. They stand in the relevant structural relations just in case the experiences (realized by neural states) in terms of which they are defined do so. And one could provide a neural or functional account of how experiences do stand in such structural relations. Finally, if felt pains are what Shoemaker calls *occurrent* response-dependent properties, Shoemaker's intentionalism accommodates Percipi.

However, one cannot answer my arguments merely by adopting response-dependent reductionism about the sensible properties. One must back this up with an account of the sensory representation (“conscious-of”) relation, which explains how we bear this relation to response-dependent properties. But, like many functionalists, Shoemaker focuses only on monadic mental *properties*, ignoring the whole difficult issue of functionalizing intentional *relations*. His only comment on sensory representation is this:

In virtue of what does an experience having a quale represent an object as having a particular [response-dependent] property? It cannot do so in virtue of a causal relation between the experience and the property it represents—one cannot say that the causing of A by B is the (or a) cause of A (here I am indebted to David Robb)... I have no fully satisfactory answer to this question (which is hardly surprising, given that no one has a fully satisfactory account of how any experience has the representational content it has) (1994, note 7).

To understand the problem, we must understand Shoemaker's physicalism. As I understand primitivism, Shoemaker counts as a primitivist. For he denies that mental properties are ‘constructions’ of more basic physical and topic neutral ingredients. In particular, he now rejects the standard functionalist claim that they are quantificational properties, defined by quantification over first-order properties. Thus Shoemaker's view on mental properties resembles Moore's view on goodness.²⁷ For Shoemaker, what goes for mental properties presumably also goes for mental relations. So Shoemaker presumably would now agree with me that the sensory representation relation is a primitive relation, not some kind of construction. Where he differs from me (and what causes the problem) is that he thinks every mental property (and presumably, relation) is realized (necessitated) by

some physical property (or relation). This is his physicalism. This commits him to accepting something I would deny (§6): that there is some dyadic physical relation between individuals and contents (property-complexes) that *realizes* the holding of the sensory representation relation between individuals and contents. When does realization happen, according to Shoemaker? Since he denies the standard functionalist claim that mental properties are role properties defined by quantifying over first-order realizers, he must reject the standard role-filler answer. Instead, he says that P realizes M just in case the forward-looking causal powers of M are a subset of those of P .²⁸ Thus Shoemaker is committed to:

- I. The sensory representation relation between individuals and intentional contents (x has an experience with content y) is the unique relation that has such-and-such forward-looking powers.
- II. There is some *physical relation* P between individuals and contents that has a *subset* of these powers, and hence (by Shoemaker's account) realizes the sensory representation relation.
- III. What externally-instantiated response-dependent properties we bear P to, and hence the sensory representation to, somehow systematically and necessarily depends on internal functional states.
- IV. P , and hence the sensory representation relation, gloms onto the very precise kind of response-dependent properties Shoemaker now thinks that it does, as opposed to the other candidates he has considered through the years.

These claims are hard to defend. I think Shoemaker is wrong to suggest at the end of the above quote that he is in the same boat as everyone else. None of the usual wide physical relations (causal relations, indicator relations) can realize the sensory representation relation, on his view. For he accepts the internalist thesis (III). In addition, as Shoemaker notes, his response-dependent properties are acausal (not to mention biologically insignificant), so we do not bear such relations to his response-dependent properties.²⁹

Nevertheless, Kriegel and Prinz have combined Shoemaker's response-dependent reductionism about the sensible properties with a Dretskean indicator theory of sensory representation, according to which a state represents what it has the function of indicating. As they focus on the general question about consciousness, the theory is underdeveloped. Prinz says that "sensations represent appearances" (2006, 451) and appearances "are powers that external things have to cause representations in our dedicated input systems" (438). But he does not say why he accepts response-dependent reductionism over response-independent reductionism, or select between the many versions of it (see the list below), or say how it meshes with his Dretskean psychosemantics. Kriegel claims that, when Norma has an

experience of blue₁₆, the determinate color appearance represented by her visual state *N* is identical with a *massive conjunction*—in his words:

The disposition to produce [neural] response *N* in Norma and subjects physically like her within her gender, race, and age group *and* to produce [silicon] response *S* in Venutian Norma and subjects physically like her in gender, race and age group *and* . . . and so on, for every type of realizer [of the experience of blue₁₆]. (Kriegel 2009, 90)

But he fails to address an obvious problem: under his favored Dretskean, teleological (or indeed any) psychosemantics, how might Norma's human visual state *N* here on Earth represents a massive conjunction involving Venutians and their genders (not to mention Martians, etc.)?³⁰

Previously, I assumed that Dretske's indicator theory of representation goes with response-*independent* reductionism about the represented sensible properties. Dretske's views agree with this assumption. Then Dretske's theory accommodates External-Directedness but not Internal-Dependence, Structure, or Percipi. Against Dretske, Kriegel and Prinz hold that Dretske's psychosemantic theory implies that the represented properties are response-*dependent*. *If the Kriegel-Prinz theory were feasible*, then it would accommodate my constraints as Shoemaker's does. In my view, this is the only case for their theory, although not one they themselves make.

But, while the Kriegel-Prinz theory is underdeveloped, we know enough to know that it is not workable. I will mention three problems.

Consider vision for a change. Suppose Maxwell views an orange. Let orange₁₇ be the "appearance" property the representation of which determines the character of his experience. In principle, orange₁₇ might be any of the following properties of the orange:

- the response-independent reflectance property *R*₁₇
- occurrently causing fine-grained neural state *N*₁₇ in Maxwell (where *N*₁₇ realizes Maxwell's fine-grained color experience)
- being disposed to cause neural state *N*₁₇ in some normal humans under some viewing conditions normal for humans
- being disposed to cause functional state *F*₁₇ (realized in humans by *N*₁₇) in the normal members of *some* species or other under viewing conditions normal for *some* species—a Shoemakerian *higher-order dispositional property*
- Kriegel's massively conjunctive property
- being disposed to cause neural state *x* in some normal pigeons under some normal conditions
- being disposed to cause reading *z* in a nearby spectrometer

It cannot be indeterminate which of these orange₁₇ is identical with. On the Kriegel-Prinz theory, it is identical with one of the

response-dependent properties to the exclusion of the others. The three problems are as follows:

Uniformity. Some properties the representation of which contributes to phenomenology are certainly not response-dependent properties of the form: *causing internal state S*. I have in mind: distances, spatial relations, locations, lengths, sizes, shapes, orientations. So the Kriegel-Prinz theory requires that our experiences represent two radically different types of properties. How might Dretske's psychosemantics deliver that result?

Incompatibility. Is Kriegel and Prinz's Dretskean (or any) psychosemantics even *compatible* with their claim that experience represents response-dependent properties of the kind listed?

Promiscuity. Supposing Incompatibility can be solved, what in Dretske's (or any) psychosemantics might determine *which* of the listed response-dependent properties Maxwell's experience represents, and hence which of them is orange₁₇? (Similar to (IV) above.)

Incompatibility and Promiscuity require development. Informally, Dretske says that, on his view, to find out what a state *s* represents we determine what property is instantiated in the environment when the system to which *s* belongs is working right.³¹ But, when a person is in Maxwell's sensory state, and optimal conditions obtain, the perceived object has *all* of the listed response-dependent properties, since they are nomically joined. Indication is a promiscuous relation. What constraints could Kriegel and Prinz add to indication in order to solve Promiscuity?

The problem is that, once we bring in the constraints on representation familiar from the psychosemantics literature in order to attempt to solve Promiscuity, we see that the Kriegel-Prinz theory founders on Incompatibility, since those constraints entail that Maxwell's experience *only* represents the response-independent property R_{17} .

The *causal constraint* is meant to whittle down candidate representata by saying that a state represents a property only if the property causes the state. (Incidentally, Dretske himself denies this.³²) But recall the Shoemaker quote. Intuitively, in this case, it is the response-independent property R_{17} , rather than any of the listed response-dependent properties, that causes Maxwell's state. So, given the causal constraint, the Kriegel-Prinz theory founders on Incompatibility. Alternatively, if Kriegel and Prinz solve the Incompatibility worry here by allowing that *some* such response-dependent properties are in some weak sense 'causally relevant', then presumably *all* of them are in that weak sense causally relevant (and hence *all* represented), so that Kriegel and Prinz still face Promiscuity.

Dretske's *teleological constraint* says that a state represents the property that it has the function of indicating. Stich persuasively argues that it is often indeterminate what a state has the "function" of indicating, so the constraint fails to ensure sufficient whittling-down.³³ Indeed, I question my grip on the notion. In any case, if Maxwell's sensory state has the function of indicating anything, it is (as Dretske himself thinks) the response-independent property R_{17} , not any of the listed response-dependent properties (especially not Kriegel's *massive conjunction*). Vision scientists say that the function of the visual system is to recover reflectances, pain researchers say that the function of the pain system is to indicate actual or potential bodily disturbances, and so on, *not* properties of the form *being disposed to cause internal state S*. So, given the teleological constraint, the Kriegel-Prinz theory founders on Incompatibility.

Responding to Stich, Dretske adds a *directness constraint*:

when an indicator, C , indicates both F and G , and its indication of G is *via* its indication of F ... then... C acquires the function of indicating that F . It is F ... that C comes to represent... Until I can develop this answer my account of representation is incomplete. (1990, 826)

Maxwell's sensory state indicates the listed response-dependent properties via indicating the response-independent reflectance property R_{17} that is their common categorical basis. So the directness constraint entails that it represents the reflectance property R_{17} , not any of the response-dependent properties. So, given this constraint, the Kriegel-Prinz theory founders on Incompatibility. Similarly for Fodor's *asymmetric dependence constraint*, since the listed response-dependent would not be indicated but for the fact that R_{17} is.³⁴

Let me address a response (thanks here to Uriah Kriegel). This response retains the claim that ordinary representation reduces to a Dretskean function-to-indicate relation. But, according to the response, Dretske's theory entails that Maxwell's state promiscuously represents *all* of the properties listed, notwithstanding all of my objections. This solves the Incompatibility problem. As for the Promiscuity problem, the response holds that it is Maxwell's representation of only *one* of these properties, namely Kriegel's massive conjunction, that determines the phenomenology of his current experience: in short, Maxwell *phenomenally* represents only the massive conjunction. This response faces many problems. I will just mention the most obvious one. On this response, what is this special "phenomenal" representation relation? It cannot be identified with the function-to-indicate relation, since on this response the function-to-indicate relation is promiscuous whereas the phenomenal representation is not. It seems that on this response the phenomenal representation relation must be a *primitive* relation, and it must be a *primitive fact* that Maxwell bears it to Kriegel's

massive conjunction as opposed to all of the other candidates. So this response requires rejecting the naturalistic Kriegel-Prinz theory now under consideration, and moving to (an implausible version of) the primitivism I favor.

The Kriegel-Prinz theory fails. Incompatibility is the main problem: their Dretskean indicator theory of sensory representation just clearly conflicts with the claim that experiences represent *any* of the above-listed kind of response-dependent properties. Elsewhere I argue that consumer-based theories do too. (I suggest that they wrongly entail that experiences only represent *other* response-dependent properties, for instance *being poisonous*.) In general, they do not help us accommodate Internal-Dependence and our other constraints. Here I disagree with Lycan.³⁵ Indeed, arguably *no* reductive psychosemantics is consistent with the response-dependent gambit. We must do better.³⁶

6. Primitivism

Only intentionalism accommodates External-Directedness (§1). And, broadly speaking, if we seek a reductive form of intentionalism, there are two possible versions. The first combines intentionalism with *response-independent reductionism* about the sensible properties. But this leads to tracking intentionalism, which violates Internal-Dependence, Structure and Percipi (§§2–4). The second combines intentionalism with *response-dependent reductionism*. But it is unavailable to the reductive intentionalist (§5). Therefore what has emerged is a systematic reason for skepticism about whether any reduction of consciousness simultaneously satisfies all of our constraints.

Elsewhere I provide another systematic argument.³⁷ Rather than rehearse the full argument for primitivism, I will explain how a primitivist version of intentionalism satisfies our constraints. On my version, both the sensible properties and the conscious-of relation are primitive.

First, *Structure*. We saw that the conjunction of intentionalism and response-independent reductionism about the sensible properties violates Structure. Response-dependent reductionism accommodates Structure but faces (*inter alia*) the Promiscuity Problem. An additional problem, not mentioned above, is its counterintuitive consequence that distinct sensible properties belonging to a common determinable (for instance, orange₁₇ and pure red₃₁) do not exclude, since they are identical with compatible response-dependent properties.

Primitivism about the sensible properties avoids these problems. It accommodates Structure because the relevant primitive properties might have the right structural features even if the correlative response-independent physical properties do not. Thus Maxwell's structure judgment [1] is not a false judgment about the response-independent physical properties D_2 and

DI, but a true judgment about primitive pain properties of which he is ostensibly conscious in his arms. Likewise for [2]–[6]. At the same time, it avoids the Promiscuity Problem, since, as we shall see, there are possible *non-reductive* explanations of why we should be conscious of certain primitive properties to the exclusion of other candidate representata. Further, it can accommodate the exclusion intuition, since the relevant primitive properties might necessarily exclude.

Once we accept primitivism, we might populate the external world with primitive sensible properties. But simplicity considerations favor *projectivism*. There exist sensible properties, and we bear the conscious-of relation to them, but nothing instantiates them, including our own experiences. They only live in the contents of our experiences.

Many scientists accept projectivism on empirical grounds. But philosophers generally dismiss it. One objection rests on common sense. However, for pain, taste and smell, *conciliatory projectivism* accommodates common sense. On this view, ‘there is a pain in *s*’s leg’, for instance, is true iff *s* is conscious of (represents) a pain property there. This can be so, even if no pain property is ever actually instantiated there. Maybe for sound and color *revisionary projectivism* is right: we naively and mistakenly attribute to objects and events the very properties of which we are directly conscious. Since an event is a sound-event only if it has sound-qualities, this would have the further consequence that strictly speaking there are no sounds: sounds, as well as sound-qualities, only exist in the contents of our experiences. This is revisionary but the common sense objection is hollow. Why should the fact that a belief is generally held *in itself* count in its favor? Another objection to projectivism is that we cannot be conscious of acausal, uninstantiated properties. But this depends on tracking theories of the conscious-of relation which we know to fail: they violate Internal-Dependence, Structure, and Percipi (§§2–4). Further, are we not intentionally directed at *other* acausal abstracta, for instance numbers? Finally, we will see that there are *non-causal* explanations of our consciousness of uninstantiated sensible properties.

Next, *Percipi*. Perhaps certain sensible properties (felt pains, tastes) are *necessarily* uninstantiated. Indeed, perhaps felt pains, tastes, and so on are not even properties. They are somewhat like properties as they can be more or less determinate and exact resemblance among them entails identity (you and I can feel the same pain). But perhaps, unlike properties, they are not *unsaturated* entities, so they are not the kind of things that can be instantiated or occur. Perhaps they are not *ways things could be*. Then they are not ways the bodily regions of an insentient cadaver could be. In general, counterexamples to the Percipi intuition [7] are impossible because its antecedent is necessarily false. This goes with the non-predicative intentionality of pain experience, taste experience, and so on: for instance, we do not feel bodily regions *as having* pains, but rather feel pains *in* bodily regions.

Finally, *Externally-Directedness* and *Internal-Dependence*. Given the following constraint on realization, simultaneously accommodating External-Directedness and Internal-Dependence appears very difficult:

Congruence: If physical property P realizes mental property M , then P must have the same ‘logical form’ as M . Thus, if the relational property *standing in the conscious-of relation to C* is realized by physical property P , then P must also be relational. In particular, it must have a relation-component that is identical with (or realizes) the conscious-of relation, as it might be the tracking relation.³⁸

Given Congruence, there must exist a dyadic physical relation between individuals and external properties that might realize, or be identical with, the conscious-of relation. The most obvious candidates are causal-indicator relations. But, as we have seen (§§2–4), if we say that the conscious-of relation is identical with (or realized by) such a relation, then, while we might accommodate External-Directedness, we fail to accommodate Internal-Dependence (and Structure and Percipi).

The trick is to reject Congruence. After all, it is not *a priori*. Then we can accommodate Internal-Dependence by saying that monadic relational properties of the form *standing in the conscious-of relation to property C at viewer relative place p* are realized by the internal physical-functional properties that differ between the individuals in coincidental variation cases. There is no need to say that the dyadic conscious-of relation that is a constituent of such relational properties is realized by a tracking relation or any physical other relation. (Likewise, standing in the thinking-of relation to the number two might be realized by a super complex internal physical-functional property that does not involve a relation to the number two.) Here are two ways to implement this approach.

Subset primitivism. This theory accepts Shoemaker’s *subset principle*: necessarily, if the forward-looking causal powers of M are a subset of those of P , then, necessarily, whatever has P has M . The neural states of individuals in coincidental variation cases have different forward-looking powers. Different properties of the form *being conscious of C at location L* have different forward-looking powers that fit those of the neural properties. By the subset principle, the different neural properties thereby realize the consciousness of different sensible properties, in accordance with Internal-Dependence. Since properties of the form *being conscious of C* can have causal powers even if the constituent property C is uninstantiated, this is compatible with projectivism.

For Shoemaker the dyadic conscious-of (sensory representation) relation must be realized by some dyadic physical relation (§5). Subset primitivism denies this. For if we accept Internal-Dependence, it is impossible see what this physical relation might be. It cannot be a tracking relation. So if physicalism is defined in terms of realization, then subset primitivism is not

physicalist. But if the subset principle is formulated in terms of metaphysical necessity, subset primitivism upholds physicalist supervenience, and so is physicalist in another sense. Elsewhere I call this *primitivist physicalism*.³⁹

Shoemaker's theory differs modally from dualism in upholding physicalist supervenience. But since Shoemaker is apparently a primitivist (§5), it resembles dualism ontologically. Indeed, it might seem worse than dualism, since the subset principle would appear to be a brute *metaphysical* necessity governing the emergence of properties. I have independently flirted with a dualist version of subset primitivism on which the subset principle is a contingent law of nature—a kind of Leibnizian principle of plenitude. The idea is that up in Plato's heaven conscious properties have certain built-in forward-looking causal powers. Complex systems evolved that have physical properties whose causal powers appropriately match those of conscious properties. By the subset principle, these physical properties bring conscious properties with them. In general, if a causally efficacious (mental or non-mental) property *can be* instantiated consistently with the laws of nature, it *is* instantiated. But there are possible Zombie worlds in which this is not so because the subset principle is false. This dualism entails rampant overdetermination. On the upside it only requires a single bridge principle rather than the usual swarm of psychophysical laws.⁴⁰

Subset primitivism is a non-reductive functionalism. Those who think functional nature cannot fully determine phenomenal nature (as it might, on the basis of black-white spectrum inversion) might prefer the second model.

Decoding primitivism. On this more biologically-oriented theory, there is, for each sense-modality, a *decoding algorithm* of the form: necessarily, if an individual is in (multiply realizable) neuro-functional state s , then he stands in the conscious-of relation to sensible property $f(s)$ at location l , where f is some fairly simple function from neural parameters onto sensible properties (and locations). The strength of the modality here might be metaphysical (yielding primitivist physicalism) or nomological (yielding dualism).⁴¹ This theory accommodates External-Directedness and Internal-Dependence. It is compatible with projectivism because it does not require that sensible properties be anywhere instantiated.

Decoding primitivism about sensory content goes against the standard functionalist view that content is always determined by links to the environment and action rather than the intrinsic features of internal states. But three arguments support it. First, on pain of mystery, there must be algorithms going from *some* physical-functional properties to sensory content. But, against the standard view, there are systematic reasons to doubt the existence of algorithms going from (input-output and internal) functional role to sensory content.⁴² The only option remaining is that there are algorithms going from neuro-functional states at a finer level of grain to sensory content. Third, since our structure judgments are causally explained by neural processing, at some level of abstraction there must be

an isomorphism between neural processing and structure judgments. Since such judgments reflect sensory content, it follows that there must be an isomorphism between sensory content and neural processing. Third, many neuroscientists believe in decoding algorithms and have provided modest evidence of their existence (§2).

However, I think this view faces an unaddressed problem concerning luck. Suppose that across-fiber pattern *A* plays the *yummy* role: it is caused by eating a certain nutritious food and causes one to get more of it. On decoding primitivism, the decoding algorithm for taste is actually *nice*: it maps *A* onto the consciousness of a yummy taste. But on this view there is a possible world in which *A* plays the yummy role as in the actual world but in which the decoding algorithm for taste is different and *not nice*: it maps *A* onto the consciousness of the actual taste of rotting flesh. In this world, on eating the nutritious food, people have a horrible taste experience, but act as if they desire it. Indeed, on a functionalist theory of desire, they do desire it. But they are worse off than us, because they desire something horrible. Why is the actual decoding algorithm for taste nice rather than not nice? Here certain functionalist theories like subset primitivism have an advantage, since they might explain why a yummy rather than yucky taste experience fills the yummy role.⁴³

In closing, let me note that primitivism is compatible with interesting theories of consciousness and its relationship to intentionality. Even if the conscious-of relation is not a physical-functional relation, it still might have internal structure. On a *first-order accessibility version* primitivism, bearing the conscious-of relation to worldly content *C* is a matter of standing in a (primitive) externally-directed relation *E* to *C* and consequently being disposed to believe *C*. On a *higher-order version*, bearing the conscious-of relation to content *C* is a matter of standing in relation *E* to *C* and standing in some internally-directed relation *I* (belief, acquaintance) to one's standing in *E* to *C*. Chalmers defends this type of view where *I* is *acquaintance*, on the grounds that it best explains privileged access.⁴⁴

I think such views are partly wrong and partly right. Against them, I think that the conscious-of relation is a wholly simple, externally-directed relation. But I think that *among believers* it is *ceteris paribus* necessarily connected with externally-directed first-order beliefs and internally-directed second-order beliefs. Phenomenology, while not analyzable (even partially) in terms of cognitive access, is necessarily linked with it. This will be so if we combine primitivism about sensory intentionality with an interpretationist-functional theory of belief and non-sensory intentionality generally.

One constraint on interpretation built into the concept of belief is that when believers have experiences they typically believe that they have those experiences. This might help explain privileged access by making second-order beliefs *a priori* necessarily super-reliable. According to Chalmers, to explain privileged access among believers we must posit an internally-directed

acquaintance relation built into all consciousness, even the consciousness of simple non-believing animals. On my view this may be unnecessary.

Other constraints on interpretation concern the relationship between experience and externally-directed first-order belief and desire, for instance that believers generally believe the contents of their experiences (recall the ‘grounding property’ from §1). This view may to *some* extent reduce Quinean worries about the determinacy of content. For, on this view, the rich primitive intentionality of experience acts as an anchor point that helps determine the contents of downstream mental states. On primitivism, not only the physical-functional facts about a person, but also the rich facts about the contents of his experiences, are part of the “given” facts that determine his beliefs and desires.⁴⁵

7. Conclusion

External-Directedness, in my view, is an undeniable feature of consciousness. It requires intentionalism. Reductive intentionalists kick qualia into the external, physical world. But decades of research in psychophysics and neuroscience show that qualia cannot be located there. Here as elsewhere reductionism has dead-ended. The way to be an intentionalist, if at all, is to adopt a projectivist and primitivist version of intentionalism. This satisfies External-Directedness, Internal-Dependence, Structure, and Percipi. I think the price in complexity is worth it. Maybe that is just the way the world is.

But those who cannot swallow primitivism might take my arguments to be a *reductio ad absurdum* of intentionalism. Empirical considerations show that, contrary to received wisdom, the best reductive theory of consciousness is not an intentional theory but an anti-intentionalist, biological theory which interiorizes consciousness and forsakes External-Directedness.

Notes

1. Prinz (forthcoming). Prinz’s theory resembles Tye’s (2006) poisedness theory; I cannot think of cases in which the theories clearly deliver contrary verdicts.
2. For references and criticism, see Loewer (1999).
3. Pautz (forthcoming b).
4. Block (2009), McLaughlin (2007).
5. McLaughlin 2007, 442.
6. Even on an “internalist” account, *N* has its content contingently, by virtue of playing a narrow functional role with respect to other internal states and behavioral outputs, which depends on the laws and other contingent factors.
7. For the property-complex version of intentionalism, see Dretske 1995, 102. On a more standard version of intentionalism, in having experiences we are in the first

instance intentionally directed at *propositions*. I work with the property-complex version for simplicity.

8. Tye (2006).
9. Dretske (1995). What tracking theories have in common is that they entail Same Experience in the coincidental variation cases introduced below. Superficially, Hill (2006) may seem to accept tracking intentionalism. But, partly in response to the sort of arguments I develop here, he is willing to adopt a “dual factor view” on which phenomenal character is in part a matter of intrinsic properties of the neural content-vehicles—the “mental paint”—though it is also, and largely, a matter of representational content (personal communication). This view rejects intentionalism in my sense.
10. Price (1999).
11. Goldstein (2007).
12. Moore (2003).
13. Moore 2003, 134.
14. Block (2006, 140) and Shoemaker (2007, 126) suggest that *actual* cases make trouble for tracking intentionalism, namely cases in which, due to individual differences, two experiences (e. g. pain or color experiences) are caused by the same stimulus. Kriegel (2009, 84–7) follows them. But elsewhere (Pautz forthcoming a, sect. 2 and sect. 4) I suggest that such a case is no problem for the tracking intentionalist. He can say that, while the actual cause of the two experiences is the same, their optimal causes (and hence the represented properties) are distinct but overlapping ranges of physical properties. Thus the phenomenal difference is grounded in a representational difference, without making either experience non-veridical. The tracking intentionalist cannot likewise reply to my hypothetical coincidental variation cases, because I stipulate complete coincidence in the physical properties optimally tracked.
15. For instance, *actaea pachypoda* (Doll’s-eyes) is a berry highly poisonous (and bitter in taste) to humans, but harmless to birds, the plant’s primary seed dispensers. Incidentally, in the hypothetical case, we might suppose that the poisonous-to-Yuck of the berries is grounded in a chemical property C^* different from C that has no causal effect on his taste-receptors, and that only three-fourths of the berries have. Hence, Yuck neither tracks nor indicates the poisonous-to-Yuck of the berries.
16. Two objections: (i) The tracking intentionalist might grant that under his theory the individuals in coincidental variation cases represent the same external physical properties, but maintain that they represent them as standing in different resemblance relations, thus accommodating Different Experiences (Byrne, in discussion). But, if sensory systems can track such relational facts (which I doubt—see the “strange relations response” in §3), then we can stipulate that the individuals in coincidental variation cases track the same such relational facts, in which case this response is inconsistent with tracking intentionalism’s psychosemantics. (ii) The tracking intentionalist might reply that in coincidental variation cases there is no difference in phenomenology at either the sensory or affective-hedonic level—this would be inconsistent with tracking intentionalism about phenomenology, since by stipulation the individuals involved optimally track exactly the same physical properties. Rather, there is only a non-phenomenal

difference in *desire* (Hill, in discussion). This response faces four problems. First, since studies show that *phenomenology* (not just desire) is better correlated with neural activity in certain regions than with any other factors (Goldstein 2007, 318–19), and since we may suppose that the individuals in coincidental variation cases differ in such activity, they likely differ in *phenomenology*. Second, the innate, fine-grained differences between the individuals in sorting and discriminatory behavior are best explained by differences not just in desire but also in innate, fine-grained phenomenology. Third, it is implausible to suppose that the individuals systematically differ in their desires but not in phenomenology, for this leaves us without an explanation of the differences in desire. Fourth, this reply in any case does not generalize to the case of Loud and Soft, or to the case of Maxwell and Mabel involving color vision that I have described elsewhere (forthcoming a).

17. The individuals in coincidental variation cases bear the optimal cause relation to the same properties but (given Internal-Dependence and intentionalism) bear the conscious-of relation to different properties. So what Internal-Dependence refutes in the first instance is tracking intentionalism's identification of the conscious-of relation with the optimal cause relation. In principle, the intentionalist could reject this claim and retain response-independent reductionism. But another, three-premise argument rules this out. To illustrate, consider the view that loudness is some super-complicated, response-independent, physical property L involving amplitude, frequency, and critical bands (O'Callaghan 2007). The first premise is that, if this view is true, then, since (on intentionalism) Loud and Soft are ostensibly conscious of (that is, represent) different loudness-levels, they must be ostensibly conscious of *different* values, L_1 and L_2 , of this super-complicated, external physical property. The second premise is that, if this is so, there must be some explanation of how it is so. The third premise is that there is no such explanation, since by stipulation what external property they bear the optimal cause relation (and kindred naturalistic relations) to is exactly the *same* (as it might be, L_1). The Structure Argument and the Percipi Argument (§§3–4) provide further reason to reject the conjunction of intentionalism and response-independent reductionism about the sensible properties.
18. One might worry that it makes no sense to say one property is twice greater than another. I disagree: one length can be twice greater than another, one mass can be twice greater than another, and so on. The properties themselves, not just the numbers we assign to them, stand in the relevant ratio relations. Indeed, an assignment of numbers to properties is valid just in case it preserves their objective ratio-structure.
19. Hill 2006, 94.
20. An additional problem with the Hill-inspired misrepresentation response to the Structure Argument is that it is not general. It does not help the tracking intentionalist accommodate certain other structure judgments, for instance 'As I view the orange, the color I'm experiencing is equally reddish and yellowish'. For suppose (implausibly) that on tracking intentionalism the represented (experienced) color is identical with a reflectance R^* that the orange does not possess. There would remain a problem: in what sense R^* is equally

- reddish and yellowish? So here misrepresentation does not help with the problem.
21. Tye and Bradley 2001, sect. 3. Hill's recent "Thouless" view of represented size likewise appeals to corrective transformations; it also faces the general psychosemantic worry to be developed.
 22. Lewis 1984. Lewis's actual theory of how expressions get their contents is more complicated than the "use-plus-eligibility" theory often attributed to him. But the details need not concern us here.
 23. There is another argument against tracking intentionalism concerning our structure judgments, one which does not even require the truth of those judgments. Good internal correlation and bad external correlation show that what structure judgments we make is explained by our neural properties, not the properties of the form *bearing the optimal cause relation to external response-independent property E* which, on tracking intentionalism, constitute phenomenology. This means that, on tracking intentionalism, when someone is making structure judgments along the lines of [1]–[6], what is driving his structure judgments about the phenomenology of his experience is not really the phenomenology of his experience, but something which, on this theory, is entirely subpersonal, namely his neural machinery. This is counterintuitive.
 24. Papineau (2002), Block (2002).
 25. McLaughlin (2007).
 26. Shoemaker (e. g. 1994) is motivated by the alleged conceivability of spectrum inversion without misrepresentation.
 27. Shoemaker 2007, 15. One might think that Shoemaker is not a primitivist since he says here that mental properties are *necessarily co-extensive with* certain quantificational properties. But, although Shoemaker's views on this matter are obscure, he still seems to be a primitivist. For he says his functionalism differs from standard functionalism. As far as I can see, the only way in which Shoemaker's functionalism might differ from standard functionalism is that it denies standard functionalism's *identification* of mental properties with the relevant quantificational properties, and with it the standard role-filler account of realization. In general, Shoemaker denies that mental properties are *identical with* any 'constructed' properties built from the stock of fundamental physical and topic-neutral properties. This makes him a primitivist in my book, even if he says mental properties are necessarily co-extensive with certain 'constructed' properties.
 28. Shoemaker 2007. I simplify greatly (for instance, I ignore backward-looking causal powers) but in ways irrelevant to what follows. Why would the causal powers of a mental property like *being in pain* be a *subset* of those of a physical property like *undergoing C-fiber stimulation*? The reason, roughly, is that C-fiber stimulation does what pain does: cause withdrawal, and so on. But it has additional causal powers that are not built into pain: for instance, it has certain local neural effects that are not built into pain.
 29. For an initially attractive but ultimately unsatisfactory solution to the problem facing Shoemaker, see the discussion of the manifestation relation in Pautz forthcoming a, sect. 7.
 30. For Kriegel's Dretskean psychosemantics, see Kriegel 2009, 66.

31. Dretske 1995, 49.
32. Dretske 1990, 819. This goes against Prinz's (2006, 441) interpretation of Dretske.
33. Stich 1990, 806.
34. For a helpful account of Fodor's theory, see Loewer 1997.
35. See Pautz forthcoming a and Lycan 2006.
36. On Klein's (2007) interesting *negative imperative theory* of pain, the phenomenology of a pain in the ankle, for instance, is fully determined by a negative imperative content along the lines of *don't put weight on it*. It might be thought that this theory helps reductionism. For instance, on this theory, pains do not represent pain qualities at all, so there is no problem of reducing them. I have three points. (i) I think the imperative theory is problematic. While negative imperatives admit of degree, it is hard to see how their degrees might match up with degrees of painfulness. What in the imperative contents of Maxwell's two consecutive pains (§3) might determine that the second pain was *roughly twice greater than the first*? (Klein tells me he addresses this worry in forthcoming work.) In addition, the theory entails, implausibly, that if one goes gradually from having a non-painful experience (e. g. of temperature) to a painful one, then the *type* of content possessed by the experience at some point radically shifts, from wholly descriptive to wholly imperative. Finally, the theory is at odds with phenomenology, since pain seems to present qualities in bodily regions. (ii) Even if it right, the imperative theory does not obviously facilitate the reduction of pain, since it is unclear how pain states in the brain might represent all of the indefinitely many fine-grained negative imperatives that are required to account for fine-grained phenomenology. Klein claims that the general function of pain is to prevent pain-causing actions. But even if this is so, it does not explain how pains *represent* specific negative imperatives. To see this, note that one function of vision is to enhance discrimination, but this does not mean that visual states represent positive imperatives along the lines of *enhance discrimination*. Further, what naturalistic facts determine that a pain experience represents a content proscribing against one specific fine-grained action rather than another? (iii) Even if it is right about pain, the imperative theory is not right about taste experience, sound experience, or color experience. So it cannot facilitate reduction for these types of experiences.
37. Pautz forthcoming a.
38. Here I use 'realizes' very broadly: '*P* realizes *M*' just means that *P* necessitates *M*. I understand realization to be compatible with dualism. On dualism, the physical necessitates the mental with nomological rather than metaphysical necessity.
39. Pautz forthcoming a, sect. 12.
40. This dualist subset theory was developed independently of Shoemaker's physicalist version and around the same time. Notice that the subset principle contains two occurrence of 'necessarily'. On the dualist subset theory I have in mind, *both* indicate mere nomic (not metaphysical) necessity, whereas on Shoemaker's physicalist subset theory *both* indicate metaphysical necessity. Of course, by the lights of Shoemaker's physicalist subset theory, the dualist subset theory is metaphysically impossible. For on Shoemaker's theory the subset principle is metaphysically necessary. So, contrary to the dualist version, there is no world in which it is nomically but not metaphysically necessary. But Shoemaker

(personal communication 2001) at least seemed to think it was *coherent*. One might think Shoemaker's physicalist subset theory has an advantage over the dualist subset theory: with a proportionality theory of causation, it avoids massive overdetermination (Shoemaker 2007, 14). But, in my view, even if Shoemaker's physicalist subset theory avoids massive overdetermination *verbally*, there remains a clear sense in which on this view higher-order properties are nomically superfluous and there is massive overdetermination. So if there is any reason to prefer Shoemaker's physicalist version to the dualist version, it does not concern mental causation. Chalmers (personal communication 2000) objected to my dualist subset theory (for which he suggested the slogan "causation loves company") that it requires (what he regards as implausible) that mental properties have built in causal powers necessarily and independently of the regularities in which they figure. But I do not think that the dualist subset theory strictly requires this claim and in any case I do not find it completely implausible. However, I do think it faces other serious problems, which I plan to develop elsewhere.

41. Decoding primitivism resembles Searle's biological naturalism (1983, 272), but without his claim that a certain fine-grained biological substrate may be necessary for consciousness. It might be wondered how my openness to decoding primitivism squares with my previous remarks on the magical view of intentionality (§1). If the decoding algorithms are metaphysically necessary, isn't decoding primitivism the same as the magical view? I have two points in reply. First, before I only said that the magical view leads to primitivism, so that it is unavailable to *the biological theorist who wants a reductive theory of the mind*. I did not say that the magical view is false. Second, if I adopted decoding primitivism, I would not say that the decoding algorithms are *metaphysically* necessary, and so would not strictly speaking advocate the magical view previously described.
42. Elsewhere (Pautz forthcoming a, especially sect. 9) I argue against the existence of a *reductive* algorithm going from (input-output and internal) functional role to sensory content. In forthcoming work I plan to provide reasons for rejecting the kind of *non-reductive* algorithm described by subset primitivism.
43. I think that certain versions of the Block-McLaughlin biological theory (§1) also face the luck argument. On this theory, the yummy experience is necessarily identical with a certain cross-fiber pattern (together perhaps with some surrounding neural machinery) *A* and the yucky experience is necessarily identical with a different such pattern *B*. We occupy a nice world in which *A* plays the yummy role and *B* plays the yucky role. But, presumably, on the biological theory, there are "not nice" worlds in which these neural states are "swapped": in such a world, we so evolved that *B* plays the yummy role and *A* plays the yucky role. (If the biological theorist denies the possibility of such a world on the grounds that *A* and *B* are "total realizers" that somehow have their functional roles necessarily, then his theory becomes equivalent to a form of functionalism.) On the biological theory, in this world, when people have *B*, they have a horrible taste experience, but act exactly as if they like it. The neural-experiential "swap" has no effect on adaptive fitness. *Prima facie*, if the biological theory is right, then the fact that we occupy a nice rather than a not-nice world cries out for explanation.
44. See Chalmers (2003). However, since 'acquaintance' is a technical term, and since Chalmers (end of sect. 4.2) suggests a deflationary interpretation on which *being*

acquainted with an experience is just *having* it, it is unclear to me exactly what the acquaintance view comes to in the end.

45. Pautz (2008). Some answer Quinean worries by saying that thoughts themselves have a primitive non-sensory, “cognitive” phenomenology which necessarily determines their (narrow) content. This is not my position; on my view, even the content of occurrent thoughts is determined by functional role. Even if there is cognitive phenomenology, the following “trumping argument” shows it cannot necessarily determine any content. Consider a possible situation in which the cognitive phenomenology *P* that actually attaches to the thought that $2 + 2 = 4$ plays the functional role in fact played by the thought that $12 + 12 = 24$. Intuitively, in that case, the cognitive phenomenology *P* does not determine that the thought has the content that $2 + 2 = 4$; the functional role trumps the cognitive phenomenology in the determination of content. (Likewise, an “absent cognitive qualia” case shows that cognitive phenomenology is not necessary for cognitive content.) In reply, it might be said that the described situation is impossible, on the grounds that a thought’s having *P* is constitutively connected with its playing the “ $2 + 2 = 4$ ” functional role; but then the view amounts to a version of the functional theory of cognitive content that I favor and does not help with the Quinean worries.

References

- Aydede, M. (ed.) 2006. *Pain: New Essays on Its Nature and the Methodology of Its Study*. Cambridge, Mass.: MIT Press.
- Block, N. forthcoming. Comparing the Major Theories of Consciousness.
- 2002. The Harder Problem of Consciousness. *The Journal of Philosophy* XCIX: 1–35.
- 2006. Bodily Sensations as an Obstacle for Representationism. In Aydede 2006.
- Chalmers, D. 2003. The Content and Epistemology of Phenomenal Belief. In Q. Smith and A. Jolic (eds.) *Consciousness*. Oxford University Press.
- Coghill, R., Sang, C., Maisog, J., and Iadarola, M. 1999. Pain Intensity Processing within the Human Brain. *Journal of Neurophysiology* 82: 1934–1943.
- Dretske, F. 1995. *Naturalizing the Mind*. Cambridge: MIT Press.
- Dretske, F. 1990. Reply to Reviewers. *Philosophy and Phenomenological Research* 50: 819–39.
- Goldstein, B. 2007. *Sensation and Perception*. Pacific Grove, CA: Wadsworth.
- Hill, C. 2006. The Paradox of Pain. In Aydede 2006.
- Klein, C. 2007. An Imperative Theory of Pain. *Journal of Philosophy* CIV: 517–532.
- Kriegel, U. 2009. *Subjective Consciousness*. Oxford: Oxford University Press.
- Lewis, D. 1984. Putnam’s Paradox. *Australasian Journal of Philosophy* 62: 221–36.
- Loewer, B. 1997. A Guide to Naturalizing Semantics. In B. Hale and C. Wright (eds.) *A Companion to the Philosophy of Language*. Oxford: Blackwell.
- Lycan, W. 2006. Pautz vs. Byrne & Tye on Externalist Intentionalism, <https://webpace.utexas.edu/arp424/www/vs.pdf>
- McLaughlin, B. 2007. Type Materialism for Phenomenal Consciousness. In M. Velmans and S. Schneider (eds.), *The Blackwell Companion to Consciousness*. Oxford: Blackwell Publishing.
- Moore, B. 2003. *An Introduction to the Psychology of Hearing*. San Diego, CA: Academic Press.
- O’Callaghan, C. 2007. *Sounds*. Oxford: Oxford University Press.
- Papineau, D. 2002. *Thinking about Consciousness*. Oxford: Oxford University Press.

- Pautz, A. 2008. The Interdependence of Phenomenology and Intentionality. *The Monist* 91: 250–72.
- forthcoming a. A Simple View of Consciousness. <https://webpace.utexas.edu/arp424/www/simple.pdf>
- forthcoming b. Why Explain Visual Experience in terms of Content? <https://webpace.utexas.edu/arp424/www/intentional.pdf>
- Porro, C., V. Cettolo, M. Francescato, P. Baraldi. 1998. Temporal and Intensity Coding of Pain in Human Cortex. *Journal of Neurophysiology* 80: 3312–3320.
- Price, D. 1999. *Psychological Mechanisms of Pain and Analgesia*: Seattle: IASP Press.
- Prinz, J. 2006. Beyond Appearances. In Hawthorne and Szabo (eds.) *Perceptual Experience*. Oxford: Oxford University Press.
- Prinz, J. forthcoming. When Does Perception Become Conscious?
- Searle, J. 1983. *Intentionality*. Cambridge: Cambridge University Press.
- Shoemaker, S. 1994. Phenomenal Character. *Nous* 28: 21–38
- 2007. *Physical Realization*. Oxford: Oxford University Press.
- Smith, D., R. Van Buskirk, J. Travers, and S. Bieber. 1983. Coding of taste stimuli by hamster brain stem neurons. *Journal of Neurophysiology* 50: 541–558.
- Stich, S. 1990. Building Belief. *Philosophy and Phenomenological Research* 50: 801–6.
- Tye, M. 2006. In Defense of Representationalism: Reply to Commentaries. In Aydede 2006.