

Can Phenomenology Determine the Content of Thought?

Peter V. Forrest

(Forthcoming in Philosophical Studies. Do not cite without permission.)

Abstract:

According to a number of popular *intentionalist* theories in philosophy of mind, phenomenology is essentially and intrinsically intentional: phenomenal properties are identical to intentional properties of a certain type, or at least, the phenomenal character of an experience necessarily fixes a type of intentional content. These views are attractive, but it is questionable whether the reasons for accepting them generalize from sensory-perceptual experience to other kinds of experience: for example, agential, moral, aesthetic, or cognitive experience. Meanwhile, a number of philosophers have argued for the existence of a proprietary phenomenology of thought, so-called *cognitive phenomenology* (CP). There are different ways of understanding the relevant sense of “proprietary,” but on one natural interpretation, phenomenology is proprietary to thought just in case enjoying an experience with that phenomenal character is inseparable from thinking an occurrent, conscious thought. While one may have instances of thought without CP experience, one will never find CP independent of thought. So the former justifiably can be said to “belong to” the latter. The purpose of this paper is to argue that these intentionalist and cognitive phenomenology views make surprisingly uncomfortable bedfellows. I contend that the combination of the two views is incompatible with our best theories of how our concepts are structured. So cognitive phenomenology cannot determine the contents of our thoughts.

Keywords: cognitive phenomenology, concepts, phenomenal intentionality, intentionalism

1. Introduction

According to a number of popular *intentionalist* theories in philosophy of mind, phenomenology is essentially and intrinsically intentional: phenomenal properties are identical to intentional properties of a certain type, or at least, the phenomenal character of

an experience necessarily fixes a type of intentional content.¹ These views are attractive.

But it is questionable whether the reasons for accepting them generalize from sensory-perceptual experience to other kinds of experience: for example, agential, moral, aesthetic, or cognitive experience.

Meanwhile, a number of philosophers have argued for the existence of a proprietary phenomenology of thought, so-called *cognitive phenomenology* (CP) (Bayne and Montague, 2011a). There are different ways of understanding the relevant sense of “proprietary,” but on one natural interpretation, phenomenology is proprietary to thought just in case enjoying an experience with that phenomenal character is inseparable from thinking an occurrent, conscious thought. While one may have instances of thought without CP experience, one will never find CP independent of thought. So the former justifiably can be said to “belong to” the latter.

The purpose of this paper is to argue that these intentionalist and cognitive phenomenology views make surprisingly uncomfortable bedfellows. This fact is surprising because many prominent advocates of CP defend it from within a broadly intentionalist framework. In particular, one might be tempted to accept the existence of proprietary cognitive phenomenology (in something like the sense just described) on the grounds that specific CP properties determine specific thought contents, such that having an experience with a specific cognitive phenomenal character suffices to entertain a specific thought. But, as I will argue, such a position is a mistake.

¹ I chiefly have in mind phenomenal intentionality (Kriegel 2013a) and other similar *non-reductive* intentionalist (Chalmers 2004, Crane 2014) theories. I focus on these theories because proponents of cognitive phenomenology tend to adopt them.

In the next section I introduce the thesis that I will be opposing: a version of intentionalism applied to cognitive phenomenology. I then elaborate on the view, suggesting a couple of subsidiary theses I think follow from the initial thesis, given reasonable assumptions. In ensuing sections, I attack this elaborated view by arguing that it is incompatible with all known viable theories of how our concepts are structured and individuated. A reasonable moral to draw from this argument is that cognitive phenomenology cannot metaphysically determine the contents of our thoughts.

2. Preliminaries

2.1 Cognitive Phenomenal Intentionality

In what follows, it will be helpful to narrow the focus of our attack onto a specific target thesis:

Cognitive Phenomenal Intentionality (C-PI): an experience with a particular cognitive phenomenal character *just is* (i.e. is type identical to) the representation of a particular thought content in a distinctively phenomenal and cognitive way.

C-PI entails impure intentionalism with respect to cognitive phenomenal properties.

Impure intentionalism is the thesis that all phenomenal properties (within a given class) are identical to certain impure intentional properties. A *pure* intentional property is the property of mentally representing a certain intentional content, and an *impure* intentional property is the property of mentally representing a particular intentional content in a particular mode or manner (Chalmers 2004). Specifying a mode or manner consists in offering a characterization of the type of mental state that is doing the representing (e.g. perceptual or cognitive, visual or auditory, belief or desire, phenomenal, conscious, occurrent, dispositional, etc), thereby telling us something of how the content is being

represented.² In the case of C-PI, the manner of representation in question is specified as cognitive and phenomenal. Note that C-PI does not commit one to the claim that the representation of *all* thought content is identical to CP experience. Assuming that cognitive phenomenology is determined by factors internal to the subject, CP thought content will be a type of narrow content.³

Although one might quibble with the details, I think most defenders of proprietary cognitive phenomenology will accept a claim similar to C-PI.⁴ For if we reject C-PI and related intentionalist theses, the idea that there is a phenomenology that is necessarily inseparable from thought is left entirely unexplained. Indeed, it is unclear why anyone would be inclined to believe it.

2.2 Presentation and Awareness

Now I want to flesh out the picture of the relationship between phenomenal character and conscious thought that emerges when we combine the C-PI thesis with certain natural assumptions. First, note that a very natural way of understanding what it is to have an experience of any kind is in terms of conscious awareness of phenomenal features.⁵ What I have in mind is summed up well in the following passage from Joe Levine:

To be phenomenally conscious, I want to say, is to be “appeared to”. It is for the conscious subject to be experientially presented with a determinate object (or

² See Chalmers 2004 and Crane 2003 and 2014, chapter 8.

³ Horgan and Tienson (2002), Siewert (2011), and Chudnoff (2015) discuss the relationship between CP contents and wide thought contents.

⁴ Some versions of intentionalism, in particular certain phenomenal intentionality views, would construe C-PI as the claim that CP properties “ground” intentional properties, rather than being identical to them. Much of what I go on to argue in this paper could be restated, with some inconvenience, in terms of grounding or some sort of dependence relation, rather than identity, or in terms that are neutral between the two. (See Kriegel 2013b; Chudnoff 2015, chapter 6.)

⁵ I understand experiences to be all and only those mental entities with phenomenology or phenomenal character (i.e. instantiating phenomenal properties).

objects) with determinate qualities...Phenomenal character, then, is not just “what it is like”, but more fundamentally, it’s “what is presented” in experience. (2011, 111)

Levine’s description of the “presentational” nature of phenomenal character seems accurate for perceptual experience. One strong consideration in favor of believing that this characterization generalizes to cases of non-perceptual experience starts from the standard gloss of phenomenology as “what-it’s-likeness.” What it is like to have an experience is always what it is like *for* a subject. The qualitative character of an experience cannot coherently be unmoored from the immediate subjective access that an experiencer has to it (cf. Burge 2007, 403ff; Crane 2014, 162-163). This fact suggests a certain apparent structure inherent to experience: there is (1) the experiencing subject, (2) an apparently two-place relation of awareness connecting subject and object, and (3) the features that constitute what it is like for the subject, i.e. the “objects” of the subject’s awareness (cf. Frey 2013, 76). To say that an “object” is “presented” to a subject in experience is just to say that the subject enjoys this direct phenomenal access to one or more phenomenal features. It does not imply that these “objects” are mind-independent, objective, public, concrete, or external entities. And because this view is committed only to the observation that experience *seems* to involve a genuine dyadic relation between a distinct subject and object, it does not imply a commitment to any particular account of the underlying metaphysical nature of phenomenal awareness, for instance whether we should prefer a sense data or direct realist theory of experience rather than an adverbialist theory. We can sum this idea up in the following thesis:

Presentation: all phenomenology consists exclusively in the presentation to the subject of apparently distinct objects (i.e. particulars and properties) in conscious experience. It is constitutive of such presentation that the subject enjoys an immediate, non-inferential awareness of these objects, which jointly constitute what it is like for her, from the subjective viewpoint, at the present moment.

Combining Presentation with our C-PI thesis yields:

Awareness: a state or episode of being aware of the objects presented in a particular cognitive phenomenal experience—objects which jointly constitute what it is like for the subject to undergo the experience—*just is* (i.e. is type identical to) the representation of a particular thought content in a distinctively phenomenal and cognitive way.

Thus, on the plausible assumption that Presentation is true, the proponent of Cognitive Phenomenal Intentionality should be led to accept Awareness as well.

2.3 Transparency

For the C-PI theorist, representing a thought content in an essentially phenomenal and cognitive manner constitutively involves (perhaps just is) a cognitive kind of conscious *awareness* of the content.⁶ For to call the representing phenomenal is to describe a phenomenally conscious episode. There might be some who would object to this step, on the grounds that what one is aware of when one consciously entertains a thought is not its content, but only various sensory-imagistic accompaniments (inner speech and the like). But presumably the proponent of C-PI would not be among these objectors, since she identifies representing the thought's content with experience of its cognitive phenomenal

⁶ Note that *representing a thought content* may fall short of *thinking a thought*. For the latter constitutively involves bearing some *attitude* to the content, and it may be that taking up the attitude involves a particular functional state, being in which is not necessitated by any phenomenal character. Not all proponents of C-PI are clear on whether propositional attitudes can be determined by CP, so I remain neutral here. Note also that although representing a thought content in a cognitive and phenomenal manner suffices for conscious awareness of that content, the latter may not suffice for the former. For one might worry about whether awareness alone can account for the "unity of the proposition": perhaps representing a complex content as a complete thought involves an act of "synthesis" of each of its components into a propositional whole, an act for which no amount of awareness of the content is sufficient. If so, representing the content is more than just awareness of it. Actually, if I am right in what I have said up to this point, then the C-PI theorist is committed to thinking that there *is* a kind of awareness that suffices for consciously representing thought contents, namely cognitive phenomenal awareness. So their view seems to commit them to thinking that synthesizing distinct conceptual contents into a proposition can be accomplished by phenomenal awareness alone (this relates to fn. 10 below). But my point is that I don't need to commit myself to this controversial position in describing their view and setting up the ensuing discussion.

features, and it is difficult to see how the latter could suffice for the former if experience of these features only allowed for awareness of sensory imagery, since merely enjoying imagery is not sufficient for representing any particular thought content.

Suppose we accept this point as well as the Awareness thesis stated above. What is the relationship between the respective *objects* of awareness in each case: the phenomenal features that constitute what it is like to undergo the CP experience, and the thought contents of which we are aware when consciously entertaining a thought? I think by far the most plausible answer to this question is some version of the following thesis:

Transparency:

- (1) The objects of cognitive awareness are introspectively *indistinguishable* from the objects of phenomenal awareness, for a given CP experience: all the properties that make up the thought content of which one is cognitively aware (perhaps along with properties that characterize the manner of one's awareness) constitute aspects of what it is like to undergo the CP experience. The two are—at least from the subjective viewpoint—one and the same.
- (2) The objects of phenomenal awareness and cognitive awareness for a given CP experience are introspectively *inseparable* from each other: one can only attend to what it's like for one to undergo the experience by attending to the thought content one is entertaining, and vice versa.⁷

The main reason proponents of C-PI should accept Transparency is that it is difficult to see how the awareness of what it's like to undergo a CP experience could necessitate, or be identical to, the awareness of what one is thinking about, unless (at least some of) *what* one is aware of is the same in both cases. It may be that Transparency is not forced upon us by this consideration. Perhaps one could allow that the awareness of specific cognitive phenomenal character metaphysically necessitates awareness of a specific thought content, but argue that awareness of the first set of objects, the qualitative properties, merely acts as a vehicle that conveys the second set of objects, the content properties, to our awareness

⁷ Cf. Crane (2014, 159-165) and Siewert (2004).

(or vice versa). This idea would involve accepting (2) in the formulation of Transparency, but denying (1): a subject is conscious of the contents of a thought *through* awareness of the qualitative features of the thought experience.

The challenge for this alternative proposal is to explain how phenomenal awareness of certain properties could necessarily mediate cognitive awareness of certain contents, despite the two objects of awareness being distinct. And this needs to be possible without the relationship between the qualitative properties and the thought contents collapsing into the merely contingent relationship between an arbitrary representation, such as a word or symbol, and what it happens to represent. If the qualitative properties presented in CP experiences acted as mere symbols for what a thought was about, they would require us to interpret them. Independent of such an interpretation, phenomenal awareness would not amount to awareness of thought content. So what would be needed for this suggestion to work is for the interpretation to be somehow “built into” the phenomenology, such that the contents of our thoughts were simply “given” to us by our awareness of the qualitative character of the experience.⁸ I am not saying that an explanation for how this could work cannot be given. But a much simpler and more introspectively plausible solution is just to claim that the what-it’s-like properties presented in CP experience require no additional interpretation to facilitate awareness of thought content, since the properties that constitute thought content also (partly) constitute what thinking the thought is like.

2.4 The Standard Account of C-PI

⁸ Cf. Siewert (1998, chapter 7) and Pitt (2011, 154, fn. 12) on perceptual experience.

In order to have a convenient label, let us call the view I have been describing over the last few sub-sections the “Standard Account” of Cognitive Phenomenal Intentionality:

The Standard Account (SA):

- Cognitive Phenomenal Intentionality (C-PI)
 - Awareness
 - Transparency

Several prominent defenders of cognitive phenomenology have advocated views that closely resemble SA. For example, in their influential introduction to cognitive phenomenology and phenomenal intentionality, Horgan and Tienson assert:

...if one contrasts wondering whether rabbits have tails with thinking that rabbits have tails, one realizes that there is something *common* phenomenologically...It is the distinctive phenomenal character of holding before one’s mind the content *rabbits have tails*, apart from the particular attitude type, be it, say, wondering, hoping, or believing. (2002, 522)

Their description of a type of phenomenology as “essentially a what-it’s-like of holding before one’s mind a specific *intentional content*” seems roughly equivalent to the suggestion that what it is like for us to think is constituted by the thought content of which we are consciously aware (541, n. 14). David Pitt (2004, 7-8), Charles Siewert (2011, 262-264), Uriah Kriegel (2015, 42-63), and Michelle Montague (2015) all express similar views.

Now I want to argue that the Standard Account is very probably false. I will do so by showing how SA conflicts with all tenable theories of conceptual structure.

3. Cognitive Phenomenology and Conceptual Structure

3.1 Cognitive phenomenal content must be conceptual content, since thoughts are composed of concepts.⁹ I will focus on *lexical* concepts: mental representations whose

⁹ I understand mental contents as *satisfaction conditions* or *conditions on extension*. This conception is neutral as to whether these contents are Russellian contents, Fregean contents, sets of possible worlds, or

contents roughly correspond to the meanings of words in natural language. These concepts are sub-propositional: they do not represent ways the world might be, but, as constituents of thoughts, they can be combined to form complex representations with propositional content. While there may be *additional* challenges that confront the Standard Account of C-PI when one considers how cognitive phenomenology could determine propositional content—how can the phenomenology be sensitive to the difference between *John loves Mary* and *Mary loves John*, for instance—thought’s compositionality means that any problems for SA we discover at the sub-propositional level should be inherited by any account of how CP relates to content at the propositional level.¹⁰

Now, when a concept is complex, in the sense that it is composed out of other concepts, philosophers and psychologists refer to its *conceptual structure*. The notion of conceptual structure includes both the relations holding between the contents of the various simpler concepts and a specification of these simpler contents, or “features” (Laurence and Margolis 1999, 4-5). A concept’s structure is what individuates the concept’s content, distinguishing it from all other conceptual contents.

The Standard Account says that a thought’s conceptual content is part of what is phenomenally presented to conscious awareness, part of what it is like for the subject to undergo a cognitive experience. So whatever the conceptual structure of a given lexical concept turns out to be, the defender of SA should be committed to the claim that everything needed to fix this structure, and thus the overall content, is presented to the

whether our mental states should be thought of as having multiple kinds of content. For convenience, I will tend to speak of the phenomenal contents of cognitive experiences as being composed of properties, rather than as modes of presentation or possible worlds.

¹⁰ For some consideration of problems that are raised for phenomenal intentionality by propositional content, see Pitt 2009, Bailey and Richards 2014, Chudnoff 2015, pp. 146-7, 151, 158.

subject in her cognitive experience. This must include (a) the *structural* (i.e. abstract relational) properties of that conceptual content as well as (b) the various individual constituents, or features, of the content.

A number of different theories of conceptual structure have been the subject of considerable debate in the literature on concepts over the last several decades.¹¹ By considering each of the main candidates below, we can see how the Standard Account for C-PI would fail given each of them.

3.2 Either (lexical) concepts have conceptual structure, or they do not. If concepts have conceptual structure, then their structure can be understood either along the lines of what Laurence and Margolis (1999) call the “containment model,” or else along what they call the “inferential model.” According to the containment model, “one concept is a structured complex of other concepts just in case it literally has those other concepts as proper parts.” On the inferential model, by contrast, “one concept is a structured complex of other concepts just in case it stands in a privileged relation to these other concepts, generally, by way of some type of inferential disposition” (5). The containment model requires, and the inferential model denies, that the other concepts out of which a complex concept is structured must be tokened every time the complex concept is tokened. We might reject these two views, and opt instead to deny that lexical concepts are structured at all. But if we accept that they do have structure, then alternatives to the containment and inferential models are not obviously forthcoming.

¹¹ See the articles collected in Margolis and Laurence 1999 for an introduction.

Arguably, the Standard Account requires the containment model of conceptual structure. For Awareness and Transparency require, in effect, that every factor that contributes to the phenomenal content of a thought is “contained” within the phenomenology of the experience. But if a concept’s structure is what specifies its semantic properties, then the structure will need to be contained in phenomenal awareness, and every conscious tokening of the concept will include its complete structure.

To appreciate this point, first consider the alternative according to which lexical concepts have no structure, but instead are atomic representations not built up out of simpler concepts or any semantically evaluable components. On the Atomistic Theory, a concept’s content consists in just its reference, which is entirely determined by the causal-historical relations the concept stands in to properties in the subject’s external environment (Fodor 1990, 1998). If Atomism were true, CP experiences would be unable to determine a concept’s content by presenting the components of its structure in phenomenal awareness, since it would have no structure. And since there is no way, simply by being instantiated by a given mental state, for CP to necessitate that that representation stands in the requisite relations to the environment that are constitutive of a concept’s content on the Atomist view, we can safely conclude that C-PI is incompatible with Conceptual Atomism.¹²

Next, consider the inferential model of conceptual structure, as typified by the *Theory-Theory* of concepts (e.g. Carey 1985, 1991). According to the Theory-Theory, the

¹² There are tricky issues here related to phenomenal externalism (Lycan 2001; Tye 1998, forthcoming), which I leave to one side, since the proponents of C-PI typically are internalists about phenomenology. For compelling reasons to reject phenomenal externalism, see Levine 2003a, 2003b and Pautz 2006, 2013. Even if we countenance the possibility of phenomenal externalism, I think Atomism will still not work with C-PI, for reasons analogous to those raised against the Theory-Theory below.

structure of a lexical concept consists in certain relations that it stands in to other concepts, which determine the inferential role it plays within this web of interconnected representations. Concepts thus resemble theoretical terms embedded in a scientific theory. Its role in the mental “theory” determines a concept’s content and identity. But then, the Theory-Theory, and the inferential model generally, are in tension with SA, which says that a concept’s content is constitutive of its phenomenal character.

For suppose a given lexical concept, such as HUSBAND, had theory structure, but also that its content was presented to a subject in phenomenal awareness. On this theory, for a concept to have the content *husband* is just for it to stand in certain inferential relations, e.g. for one to be disposed to infer MALE and MARRIED from it. But now, first, what exactly is the subject aware of when she is phenomenally aware of *husband*? By stipulation, she cannot be aware of the constituent contents *male* and *married*, at least not without this view collapsing into a version of the containment model. But then how is she aware of the property of husband-hood, under what “guise” is this the object of her awareness? What is the essence, the individuating core features, of being a husband, if it cannot involve any reference to being a male or being married? Presumably, in response one should say that the CP character need not “contain” the contents *male* and *married* as objects of awareness in this way. Rather, this phenomenal character gets to count as the content, *husband*, just in case *whatever* one is phenomenally aware of when one is entertaining this concept determines that one is disposed to draw those inferences that constitute its theory structure.

But, second, *how* is this phenomenal character supposed to determine that the concept stands in the requisite inferential relations to other, distinct and independent concepts? What about the phenomenal features of the CP experience of *husband* makes it the case that one will be disposed to infer *married* simply by undergoing this experience? The idea is entirely magical. Entirely magical, that is, unless awareness of its conceptual components is somehow inherent in the awareness of *husband*. For it is hard to imagine what *other* aspect of this content, immediately presented in experience, could be such that attending to it necessarily disposed a subject to make the requisite inferences, unless it just boiled down to the subject “seeing” that husbands are married males. But that would be to fall back on the containment model again. So either CP cannot fix inferential structure at all, or else it only does so because concepts have containment structure.

Of course, I am not denying that entertaining a concept could dispose one to entertain other concepts without one’s entertainment of the first concept needing to “contain” entertainment of all the others. There could be causal-functional facts about the representational *vehicles* of the concepts HUSBAND and MARRIED, for instance, that ensure the right inferential relations hold; the subject is just cognitively set up in such a way that when she deploys the first concept she is thereby disposed to deploy the second. And these facts would normally suffice for the Theory theorist to explain how concepts have the structure they have. But the Theory theorist who is also a C-PI theorist cannot appeal to these facts alone to explain the inferential structure of these concepts, since if she did, we could always ask her why the particular CP characters for these concepts couldn’t be separated from facts about their conceptual vehicles, and it seems she would have no

answer. Plainly, until the proponent of this view can offer the *rudiments* of an explanation as to how CP character could metaphysically determine inferential structure, we are better off considering a different version of conceptual structure that fits more naturally with SA.

This line of reasoning leaves us with the containment model. I will present two arguments for the conclusion that the Standard Account and the containment model are incompatible: the problem of Phenomenal Explosion and the problem of Structural Inadequacy.

4. The Problem of Phenomenal Explosion

4.1 Two of the standard theories of conceptual structure, the *Classical or Definitional Theory*, and the *Prototype Theory*, are most naturally understood as adopting the containment model (Laurence & Margolis 1999).¹³ For simplicity's sake, let us focus on definitional structure. According to the Definitional Theory, a concept's structure constitutes its definition: necessary and sufficient conditions for the correct semantic application of that concept, i.e. for something to count as falling within its extension. Since on the containment model of the Definitional Theory, a concept is literally built up out of its definitional structure, this structure is contained within any individual token of the concept.

For instance, suppose being a male and being married are parts of the definitional structure of the concept HUSBAND. Then, since employing the concept must involve these two components—categorizing something as a husband quite literally involves categorizing it as male and as married—HUSBAND cannot be tokened without also tokening

¹³ For an example of the former, see Peacocke 1992, and of the latter, see Rosch and Mervis 1975.

its constituents MALE and MARRIAGE. Furthermore, according to the Definitional view, MALE and MARRIAGE presumably are also concepts with structure, and decompose into other concepts that constitute necessary and sufficient conditions for their application, until at some point this decomposition terminates in primitive, unstructured concepts.

Combined with C-PI, this theory entails that the CP experience that represents the conceptual content *husband* includes component experiences that represent *male* and *marriage*, which in turn include components that represent the definitions of these contents, and so on. The inevitable result of this combination of views is that for a great many everyday concepts, entertaining them requires undergoing experiences with a dizzying degree of phenomenal complexity. Even a very brief and simple thought, such as *the husband is late to the birthday party*, will involve undergoing a combination of several of these extremely complex CP experiences. This result, I submit, is psychologically and introspectively dubious, to say the least.

The problem is not only that the entire definitional structure of a given concept would have to be *determined* by some aspect of cognitive phenomenology every time we entertain a thought with that content. That much is guaranteed just by C-PI. But once we add the Awareness and Transparency theses, the view also requires that the subject can be *aware* of all those properties contained within the concept's definition. So the challenge is having an absurdly large number of properties, and their relations, *presented* at once by the phenomenal character of a single experience. Call this the problem of *Phenomenal Explosion*.

4.2 The defender of the Standard Account might object at this point that the prospects of Phenomenal Explosion do not pose much of a difficulty for SA at all, on the grounds that *perceptual* experiences—for example the visual scene of the lit room before you at this moment—involve a great deal of phenomenal-intentional variety and complexity, presented to our conscious awareness at every waking moment. Since this degree of richness in what it's like for the subject is not only possible but also manifestly *actual*, there is no reason to think something similar could not occur in the case of conscious thought.

An initial reply to this point is that our experience of our thoughts does not seem nearly as rich and various as a typical perceptual state (does your split-second awareness of *breakfast* in the thought *what did I have for breakfast this morning?* seem to involve anything like the level of detail of your visual experience of a crowded room?). So if the Standard Account predicts that a CP experience would involve anything like the complexity of a perceptual experience, then that already is a problem.

But what's more, I believe SA demands that cognitive phenomenal awareness contains levels of complexity not required of perceptual awareness. To appreciate this point, we can compare SA with an analogous intentionalist account for perceptual experiences.

Consider the concept, ORCHESTRA, the definition of which, suppose, includes the notions of *violinist*, *cellist*, *conductor*, and so on. According to the Standard Account, whatever sets an orchestra apart from a band or choir must be phenomenally presented in the CP experience for *orchestra*, or else there would be nothing to make it the case that the subject was thinking about one rather than the other. So the overall experience must include phenomenal features corresponding to *violinist*, *cellist*, *conductor*, and so on, as well

as the relations between these elements constitutive of being an orchestra. Nor are the relations that need to be captured by the experience limited to relations that explicitly form part of the content: the relation of playing-a-different-instrument-than, which holds between violinists and cellists, is not explicitly part of the definition of ORCHESTRA. But this relation still must be *retrievable* from the phenomenology for this concept, in the sense that, at least in principle, the (suitably idealized) subject must be able to “read off” this relation from the noticeable features of the content that *are* presented by the experience, since it is entailed by those features. By analogy: my visual experience of a group of basketball players includes the presentation of their respective heights. But the phenomenology of the experience also captures the fact that player A stands in the taller-than relation to player B even if this property is not explicitly presented by the phenomenology.

In short, the Standard Account requires that there be a structure-preserving mapping, or *isomorphism*, between the properties represented in a concept’s content and the qualitative properties of that concept’s CP experience. Every relation (that makes for genuine similarity and difference) that holds essentially between the properties represented in the content maps onto either an introspectively discernible relation holding between the qualitative properties, or else a relation holding between them that could be “read off” (i.e. inferred by a suitably idealized subject, based on attending to) those properties.¹⁴

¹⁴ Isomorphism: Two sets, A and B, are isomorphic if and only if there is a one-to-one map between them, and there is some relation defined over the members of A and some relation defined over the members of B such that members of A stand in the former relation if and only if the corresponding members of B stand in the latter relation (Kulvicki 2014, 196).

Now, in the case of *perceptual* experience it seems intuitively plausible that there always will be this isomorphism between phenomenal properties and represented features of the subject's environment. For example, if I have an experience of pain in my chest composed of three consecutive qualitative properties, the content of this experience will posit three distinct pain sensations rather than two or four. And these sensations will be represented as being instantiated in space relative to other felt bodily sensations in a way that preserves the structure of the relationships that the various sensory qualities are experienced as standing in to each other. If the second pain quality seems more like the first than the third—e.g. one and two feel more stabbing, while three feels more dull—then it will be part of the content of the experience that the second pain is more like the first than the third.

However, the challenge of preserving the relational structure of the properties represented by a perceptual experience within its phenomenology is eased considerably by the fact that those properties of our environment that we perceive naturally form *families*, the members of which all differ from one another to varying degrees in a handful of well-defined ways.

As David Rosenthal (1991, 2006) has argued, families of sensory phenomenal properties are necessarily *homomorphic* to the families of perceptible properties: “the members of these families resemble and differ from one another in ways that parallel the similarities and differences among the corresponding perceptible properties of physical objects” (Rosenthal, 1991, 22-23).¹⁵ The members of each family of perceptible properties

¹⁵ A *homomorphism* is any (partial) structure-preserving map between two sets of things, of which an isomorphism is a more demanding type. The difference between the two does not matter for our purposes.

can be arranged in a multi-dimensional quality space according to a similarity metric defined by objective facts about which physical properties the relevant human perceptual modality can discriminate between, in what ways, and to what degrees (see Rosenthal 2006, 196-203). The upshot is that the properties that constitute the contents of our perceptual experiences are ideally suited to be presented in conscious awareness as complexes of sensory qualities, since, in virtue of what they're like for the subject, sensory qualities also group naturally into families.

In contrast, most of the worldly properties for which we have lexical concepts do *not* intuitively fall into mutually exclusive families of properties, the members of which all relate to each other in only a handful of well-defined ways (relevant to the individuation of each concept), which can ground relations of relative similarity and difference between the members and allow each of them to be assigned a specific location in a quality space. As Rosenthal has noted,

Homomorphisms cannot operate on the totality of everything one can form a concept of, since there are no well-defined similarities and differences that hold among all such things...Dogs bear no suitable relations of similarity and difference to everything else we can form a concept of; so we can't appeal to such relations to fix the concept of a dog. (2006, 208)

Let me be clear: unlike Rosenthal, I am not proposing that we attempt to *fix* the content of our concepts and thoughts purely by appeal to the various relations that these contents stand in to each other. My point is that the proponent of SA still needs to explain how the structure of a concept's content, *including the relations that its constituent features stand in to one another*, can be *matched* by what is presented to the subject in phenomenal awareness. The relations between, e.g., all the possible sounds made by the various instruments in an orchestra can be matched using phenomenal properties that differ along

just three qualitative dimensions (pitch, loudness, timbre). However, the relations between the types of musical instruments themselves cannot be mirrored by qualitative differences ordered along only three dimensions. Intuitively, we might say a violin is more like a viola than like a trumpet; but there is no straightforward answer to the question of whether a violin is more or less like a trumpet than like a saxophone. In contrast to the perceptual representation of colors or sounds in the room before you, the entire definitional structure for a complex concept such as ORCHESTRA in all likelihood could not be captured in awareness by only a handful of phenomenal primitives differing along only a handful of qualitative dimensions.

Thus, not only must the CP for a single thought contain an explosion of distinct phenomenal types—as would a visual experience of many distinct shades of blue, for instance—but also an explosion of different ways these phenomenal features are discernibly related to each other. And all this in the time it takes for a thought to flash through consciousness. I am not sure this is strictly *impossible* for cognitive phenomenology. But I am saying the idea that we have these experiences every time we consciously entertain a thought is just not credible. And given the limitations on consciousness in the human mind, I doubt normal humans could ever have such experiences.

4.3 I want to take a moment to address a couple of objections to the preceding argument. First, one might worry that my reasoning depends on the assumption that a subject can only be aware of a whole if she is aware of all of its parts. But here is a

counterexample to this principle: I see a brick wall from a distance; I'm aware of the wall, but not aware of each brick.¹⁶

In fact, my argument assumes no such principle. To see why, let's take a closer look at the brick wall example. In the example, am I supposed to be visually aware of the brick wall *as* a brick wall, i.e. am I aware of the property of being a brick wall? If not, then this sort of awareness doesn't provide a good analogy with CP awareness. For according to the Standard Account, phenomenal character is supposed to make one aware of, not just the object of a particular thought, but also the various properties by which the thought characterizes that object. But if my brick wall visual experience presents me with the wall but not as a brick wall, then it seems the content of the experience is consistent with the wall being something else, say a stone wall or a cardboard copy of a brick wall. And then the phenomenology can't be said to fix the content to be about brick-wall-ness.

On the other hand, suppose I am visually aware of the brick wall as a brick wall. Then of course I can be so without being aware of every brick (perhaps of *any* particular brick). But what I cannot do is be visually aware (in the experiential sense fixed by visual phenomenology) of its being a brick wall without being visually aware of its being made of bricks. This is what's analogous to the idea of being aware of a conceptual content without thus being aware of the essential, individuating features of that content. Without providing awareness of these features, how does the phenomenal character of the experience necessarily succeed in singling out the right concept to employ? But if it cannot do that, then it seems the CP of the thought cannot determine its content.

¹⁶ I am indebted to an anonymous reviewer for both the objection and the example.

Moving on, someone might suggest that we can avoid Phenomenal Explosion by appealing to the “division of cognitive labor”: perhaps the cognitive phenomenology of *my* thoughts about, e.g., forms of political rule do not need to contain the entire definition of aristocracy, nor even contain detail sufficient to distinguish *aristocracy* from *oligarchy* or *plutocracy*. What matters is just that the political theorist’s CP experiences distinguish between these contents, and that my CP experiences defer somehow to the experts.

I have two responses. First, the experts’ thoughts would still face Phenomenal Explosion, which as a description of how anyone thinks is hugely problematic (and a great many people have expert knowledge of *some* domains). Second, we can concede that the conceptual structure of many of our concepts does not amount to full-fledged definitions, but the problem of Phenomenal Explosion does not arise only for the Classical theory of concepts. The same issue must be faced if our CP experiences contain partial definitions, prototypes, exemplars, or indeed any kind of structure that is at least somewhat complex. The problem could be avoided if the majority of our concepts consisted simply in an appeal to “whatever the experts say about x”. But if this kind of passing the buck were all that conceptual content amounted to, then this would defeat the purpose both of positing conceptual structure in the first place, and of the idea that conscious thought involves “holding before one’s mind” specific thought contents. For surely no C-PI theorist thinks that our awareness of our concepts is exhausted by a string of introspectively indistinguishable appeals to what the experts think about x, y, and z.

The last paragraph illustrates why Phenomenal Explosion is not just a problem posed by the wideness of our thought contents: even if we assume that concepts commonly thought to have wide contents, such as social and natural kind concepts, will have narrow

correlates that are shared by phenomenal duplicates, these narrow contents presumably will still have complex internal structures that must be captured in their cognitive phenomenology.

This point marks a crucial difference between the Phenomenal Explosion argument and an important argument recently given by Elijah Chudnoff (2015, chapter 6). Chudnoff argues as follows. (1) The contents of many of our concepts—including demonstrative, indexical, nominal, natural, mathematical, artifactual, social, normative, and logical concepts—constitutively depend on one or another of the following factors: the objects in the subject's external environment; the non-superficial underlying physical and mathematical natures of things; the subject's social environment; and the subject's inferential dispositions. (2) These factors plausibly do not supervene on phenomenology alone: two individuals might be identical with respect to phenomenology but differ with respect to any of these factors. (3) Every one of our thoughts involves at least one of the types of concept on the above list. So, (4) cognitive phenomenology by itself cannot determine the contents of our thoughts.

The natural response on behalf of C-PI is to contend that there is a layer of conceptual content, distinct from that described by this argument, that is determined by phenomenology alone. On the plausible assumption that phenomenology is narrow (fixed by facts "in the head"), the content in question must be a (particularly demanding) kind of narrow content. So Chudnoff's response to this response consists of an extended attack on various theories of narrow content.

Chudnoff no doubt would be the first to admit that his arguments in support of his first premise, and those attacking the proposal of a distinct layer of narrow thought

content, are largely standard, well-known considerations in the literature on content internalism/externalism. This means that for the most part the worries he raises for phenomenally determined thought content are not specific to the debate about cognitive phenomenology, but rather general worries for any form of narrow content. As a result, they are unlikely to bother those philosophers, like many phenomenal intentionality theorists, who already are happily committed to narrow content. By contrast, my argument presents a specific difficulty for C-PI that is independent from any skepticism about content internalism in general (indeed, even from skepticism of the idea that content can be determined by phenomenal character). To see this point, note that even if we deny the second premise of Chudnoff's argument—in favor of the view that environmental and dispositional factors help fix thought content, but are *themselves* fixed by what's presented in cognitive phenomenal awareness—we will still face the problem of Phenomenal Explosion. But for Chudnoff, the issue just seems to be whether or not we acknowledge that his list of non-phenomenal factors plays a role in fixing content—as illustrated by the fact that he takes denying his second premise, e.g. by accepting phenomenal externalism, as a way to save C-PI (148-149).

The notable exception in his arguments comes on p. 158, where he appeals to his earlier conclusion that logical and normative concepts constitutively depend on inferential dispositions (146-147) to suggest that such dispositions might play an essential role in fixing the narrow contents of our thoughts in general. If so, then even though the content in question would still be narrow—since inferential/conceptual roles are in the head—it wouldn't be determined by cognitive phenomenology, since the earlier case of logical concepts shows that dispositions don't supervene on phenomenology. What's distinctive

about this point is that it attacks, not narrowness in general, but the notion of content that depends on CP.

The problem is that Chudnoff does not provide an argument for the claim that in general the narrow contents of concepts depend on inferential and belief-forming dispositions, nor for thinking that dispositions in general couldn't supervene on CP, beyond appealing to his one logic example. The defender of C-PI could respond by simply rejecting one of these claims: either inferential dispositions don't play a role in fixing most conceptual contents, or to the degree that they do they are themselves fixed by CP.¹⁷ We want to know why this response is unsatisfactory.

By contrast, I offer an argument for thinking that this response is unsatisfactory. Recall that in the discussion of the inferential model in section 3.2 I argued that if C-PI is true, then concepts could only have inferential structure if they *also* have containment structure. But then, commitment to the combination of C-PI and the containment model leads to Phenomenal Explosion. Likewise, if we eschew appeal to inferential structure entirely, in favor of the view that all narrow contents are determined by what is "internal" to CP awareness alone, we face Phenomenal Explosion. So either way, we need to appeal to some factors—inferential role, environment, etc—"external" to cognitive phenomenology. However, while for Chudnoff such a claim is a *premise* in his argument against C-PI, in my approach it's a *moral* to draw based on our reasoning, a pay-off to which we're entitled only at the end of the argument.

¹⁷ Which of these two options is more attractive may depend on whether the subject in question is suitably idealized. In an idealized reasoner, it is not crazy to think CP experiences that presented a concept's structure could ground meaning-constitutive dispositions to use that concept. With ordinary subjects, the better option for C-PI may be to deny conceptual contents depend on having the right dispositions.

5. The Problem of Structural Inadequacy

Let's return to the main thread of the argument. Even if—contrary to the above—we had complex CP experiences for every concept, the constituents of which were isomorphic to that concept's definitional structure, these experiences could still fail to fulfill Awareness and Transparency. This point leads us to our second argument that the Standard Account is incompatible with a containment model of conceptual structure: an experience whose cognitive phenomenology has the same *relational* structure as a concept's content is not, for that reason alone, presenting the content in phenomenal awareness. For the identity of the content also depends on the identity of its basic, non-relational elements. And SA will have an insuperable problem accounting for how the various primitive features of a concept's content can be presented in phenomenal awareness. Call this the problem of *Structural Inadequacy*.

To see the problem, consider by way of contrast the theoretical options available to a containment model of concepts that does *not* involve cognitive phenomenology, such as a standard Definitional Theory. A non-phenomenal Definitional Theory will explain how the majority of concepts fix their contents by appealing to their definitions. But definitions need to end somewhere: in order to account for the semantic contents of the primitive features that combine to make up a definition, the definitional theorist is free to appeal to a different theory, such as Conceptual Atomism or the Theory-Theory. So on this account, our conceptual repertoire is built up by containment structure out of a relatively modest handful of (innate or learned) primitive concepts that acquire their content some other way, e.g. via causal-teleological links to the environment. But this escape route is not open to the C-PI theorist since, as we have seen, these alternative theories of conceptual

structure are not compatible with the Standard Account's central idea that everything necessary to determine a concept's content must be contained in what is phenomenally presented to the subject.

For example, on a non-phenomenal definitional account the definition of the concept ARISTOCRACY might involve something like the following content: *a form of government in which political power is held by the nobility*. The content *nobility* might include (*inter alia*) the content *hereditary ruling class*, which might include the content *parent*, which might include the content *animal*. But ANIMAL, the definitional theorist may claim, not implausibly, is a primitive concept. And she can maintain that primitive concepts are atomic: ANIMAL has no structure, but gains its meaning from being reliably hooked up to the subject's environment in the right sort of way. But what should the defender of C-PI say about a primitive concept such as ANIMAL? It clearly has a rich meaning that makes an important contribution to the overall content of ARISTOCRACY and thus must show up in phenomenal awareness—but how? I see three main options for the C-PI theorist here, none of which look promising.

Option 1: Simple Qualities

Perhaps conceptual features are presented to a subject in experience as pure, simple cognitive-phenomenal qualities. Up to this point I have been using 'qualitative' as synonymous with 'phenomenal'. But here, by calling a property a "quality" I mean our grasp of its intrinsic nature—in the present case, our grasp of the ineffable, subjective character that constitutes what an experience is like—is conceptually independent from structural and functional concepts. This independence is often thought to lead to an "explanatory gap" between our understanding of the quality and our understanding of the

properties that form the basis of both everyday and scientific explanation (Levine 1983, Chalmers 1996).

On the present suggestion, all the phenomenal primitives out of which CP experiences are constructed are simple qualities similar to colors or smells. These qualities cannot point us towards non-qualitative physical-functional content in isolation, but only as they are presented in complex configurations and stand in discernible relations to one another. For by stipulation, their intrinsic natures do not contribute to conveying information about the physical world, except insofar as they are discernibly more or less similar to other qualities. By contrast, as the case of *animal* illustrates, the conceptual primitives themselves are *not* likely to be completely simple. Most of them will have an intrinsic nature with some degree of internal complexity and variation, which an utterly simple quality cannot capture. As a result, the burden of determining thought content that this view places on the relations holding between primitives is too much for them to bear. Mere relational structure will leave something essential out, namely the contents contributed by the conceptual primitives themselves.

Option 2: Just More Structure

A natural response to this challenge is that qualities *could* be sufficient to present conceptual features to the subject, provided these features are much simpler than we were just supposing. Our mistake was in thinking that a lexical concept with complex content such as ANIMAL could be a primitive feature of ARISTOCRACY, when in fact it must decompose further into much simpler concepts. In other words, what we need is to introduce far more conceptual structure before we reach conceptual bedrock.

After all, for all we know the world according to our best physics could turn out to be structure all the way down. We cannot assume that a sufficiently detailed structural description of the microphysical facts that underlie a concept's extension necessarily would fail to distinguish that concept's content from all other contents. On the contrary, the conceptual primitives could appear in experience as simple qualities, for provided they are combined and related to each other in sufficiently intricate and varied ways, they should be able to capture the individuating "essences" of even the most sophisticated conceptual contents.

But obviously most of our concepts are not that theoretically sophisticated. The structures of our ordinary concepts do not reach down that far into the underlying nature of things. CP character does not make us phenomenally aware of all of the world's structure, right down to the building blocks of reality at the quantum level, every time we entertain an ordinary, macro-level concept! If the Standard Account required awareness of abstract relational structure with even a fraction of this level of detail and sophistication, then the price of phenomenally individuating our concepts would be to introduce an even more absurd degree of complexity into each CP experience than we considered earlier.

At this point in the dialectic, Structural Inadequacy does not so much pose a problem distinct from Phenomenal Explosion as it illustrates just how serious a problem the latter can become. The initial problem of Phenomenal Explosion was that our cognitive experiences intuitively do not present the structural complexity that our concepts might plausibly be thought to possess. But on the present proposal, the problem is that in order to individuate our concepts our cognitive experiences would have to present a level of

complexity in containment structure beyond what anyone would think our concepts possess in the first place.

Option 3: Intentionally Rich Primitives

It is debatable whether all phenomenal properties really are qualitative, if that entails, e.g., that they can be used in standard anti-materialist arguments such as the Knowledge Argument.¹⁸ Perhaps some experiences just amount to being immediately aware of ordinary properties and things, which can be fully described in commonsense and/or scientific terms without any mysterious, irreducibly qualitative remainder. If so, then we might suppose that the primitive contents of lexical concepts are not presented in experience as pure qualities, but as the complex physical, functional, or structural properties that they are. In other words, perhaps the cognitive phenomenal atoms are themselves intentionally rich. This would explain how a conceptual content like *animal* could be a primitive feature.

If we adopt this proposal, we might think that the problem of Structural Inadequacy dissolves. Experiencing a concept still will involve being aware of all the properties and relations that contribute to its structure. But the content of a CP experience will not have to be derived entirely from the relations holding between purely qualitative building blocks, since awareness of conceptual features is intrinsic to the phenomenal ground floor, so to speak.

But this proposal looks to me more like offering a description of the problem than like providing a solution to it: to say that the phenomenal primitives must be richly intentional is not to explain how this could be achieved. Presumably the solution must

¹⁸ On this issue, see Bayne and Montague 2011b, Carruthers and Veillet 2011, and Thompson 2010.

involve the primitives possessing a certain degree of internal variation and complexity. It is hard to see how one could be phenomenally presented with the content, *animal*, without that involving immediate awareness of at least some details about what it is to be an animal. So an intentionally rich CP primitive must be built up out of discernibly distinct phenomenal elements, even if the primitive experience is a gestalt, in the sense that these simpler elements cannot be pulled apart and experienced independently of the phenomenal whole.

But how complex must these primitives be? They would have to be simple enough to avoid the threat of further phenomenal explosion, but complex enough to make it plausible that by combining a reasonably small number of them a single CP experience could succeed in individuating a lexical concept. The problem is that it is extremely difficult to think of any real-world properties that might meet these two criteria. The only decent candidates for the role seem to be the sensible properties represented by low-level sensory-perceptual experiences: colors, shapes, spatial locations, motion, sounds, smells, temperature, pressure, texture, and so on (perhaps along with a few others, such as objecthood and causation). We know from perceptual experience that these properties are simple enough to be phenomenally presented, and arguably many everyday concepts have a perceptual basis, being constructed from stored experiences of these properties.¹⁹

Unfortunately, the empiricist tradition, which sought to base all human concepts on our ideas of sensible properties, does not have a particularly happy history.²⁰ I doubt very much that proponents of C-PI would be happy to accept a full-blooded empiricism about

¹⁹ For a modern defense of this historically important and intuitive idea, see Prinz 2002.

²⁰ See Margolis and Laurence 2014, section 3.2 to get a flavor of the serious difficulties that confront concept empiricism.

concepts just to save their theory. Admittedly, I have not offered an *argument* for thinking that only sensible properties are remotely suited to play the role of basic conceptual features. So my point should be formulated as a challenge to defenders of C-PI: give us a realistically detailed alternative account of what these conceptual building blocks—which can be captured by intentionally rich cognitive phenomenal primitives—could consist of, or else be stuck with concept empiricism.

To sum up, the problem that Structural Inadequacy raises for the Standard Account is that mirroring the abstract structure of a concept's content in cognitive phenomenal character is insufficient to present that content to the subject, unless the primitive conceptual features are also presented. But we have seen no way to present these features in experience that does not either greatly exacerbate the problem of Phenomenal Explosion or else rely on a highly dubious empiricism, according to which CP experiences are constructed out of experiences of sensible properties. This argument, in conjunction with the argument from Phenomenal Explosion, establishes that CP experiences are inadequate vehicles for carrying a lexical concept's complete (narrow) content.²¹

6. Conclusion

This completes my argument for why our thought contents cannot be fixed in the way described by the Standard Account of C-PI. We might want to resist this line of reasoning, on the grounds that, for all we know, the *true* theory of conceptual structure

²¹ It should be clear that other versions of the containment model will not help here. For instance, one might think to turn to Prototype Theory, which says that conceptual structure consists in a statistical analysis of what properties the things that fall under a concept are likely to have (Laurence and Margolis 1999, 27ff). If prototype structure consists in complete sets of prototypical features, *contained* in the concepts themselves, then this theory cannot avoid the problems that confront the definitional view. Similar remarks apply to Exemplar Theory (Medin and Schaffer 1978).

may vindicate the Standard Account. Maybe the problem is not with C-PI, but with our current best theories of human concepts. But the success of this response is entirely hostage to future developments in cognitive science for which we have no evidence. And I think it is very unlikely that theories of conceptual structure *will* develop in a way that is congenial to the Standard Account, since the reasons for their incompatibility adduced here seem to apply to any theories of concepts that fall into the three broad categories we considered: the containment model, the inferential model, and the atomist alternative. Indeed, it is difficult to see how any concept *could* have internal structure that was not related to it either by containment or by inference.

The problem, rather, lies with C-PI's model of thinking as a form of experiencing, where this is naturally understood to involve a subject's conscious awareness of a thought content, "holding it before her mind." Once we allow that thoughts have an internal conceptual structure (and if they don't, what precisely is being held before the mind?), it becomes clear that on any plausible theory of what our thoughts are about, there will simply be *far too much structure* for us to capture it in normal human conscious experience. So the prospects of integrating cognitive phenomenology with intentionalism, in anything like the way described by the C-PI thesis, look dim.

References

Bailey, Andrew R. & Richards, Bradley. "Horgan and Tienson on phenomenology and intentionality." *Philosophical Studies*, 2014: 167 (2):313-326.

Bayne, Tim, and Michelle Montague, . *Cognitive Phenomenology*. Oxford: Oxford University Press, 2011a.

—. "Cognitive Phenomenology: An Introduction." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 27-28. Oxford: Oxford University Press, 2011b.

Burge, Tyler. "Reflections on Two Kinds of Consciousness." In *Foundations of Mind*. Oxford: Oxford University Press, 2007.

Carey, Susan. *Conceptual Change in Childhood*. MIT Press, 1985.

—. "Knowledge acquisition: Enrichment or conceptual change." In *Concepts: Core Readings*, Eric Margolis & Stephen Laurence (eds.), 459--487. MIT Press, 1999.

Carruthers, Peter, and Benedicte Veillet. "The Case Against Cognitive Phenomenology." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 35-56. Oxford: Oxford University Press, 2011.

Chalmers, David. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press, 1996.

—. "The Representational Character of Experience." In *The Future for Philosophy*, edited by Brian Leiter. Oxford: Oxford University Press, 2004.

Chudnoff, Elijah. *Cognitive Phenomenology*. Routledge, 2015.

Crane, Tim. *Aspects of Psychologism*. Harvard University Press, 2014.

—. "The Intentional Structure of Consciousness." In *Consciousness: New Philosophical Perspectives*, by A Jokic and Q Smith (eds). Oxford and New York: Oxford University Press, 2003.

Fodor, Jerry A. *A Theory of Content and Other Essays*. MIT Press, 1990.

—. *Concepts: Where Cognitive Science Went Wrong*. Oxford University Press, 1998.

Frey, Christopher. "Phenomenal Presence." In *Phenomenal Intentionality: New Essays*. Oxford: Oxford University Press, 2013.

Harman, Gilbert. "The Intrinsic Quality of Experience." Edited by James Tomberlin (ed.). *Philosophical Perspectives 4, Action Theory and Philosophy of Mind* (Ridgeview), 1990.

Horgan, Terence, and John Tienson. "The Intentionality of Phenomenology and the Phenomenology of Intentionality." In *Philosophy of Mind: Contemporary Readings*, edited by David J. Chalmers. Oxford: Oxford University Press, 2002.

Kriegel, Uriah, ed. *Phenomenal Intentionality*. Oxford University Press, 2013a.

—. "The Phenomenal Intentionality Research Program." In *Phenomenal Intentionality: New Essays*, edited by Uriah Kriegel. Oxford: Oxford University Press, 2013b.

—. *The Varieties of Consciousness*. Oxford: Oxford University Press, 2015.

Kulvicki, John V. *Images*. New York: Routledge, 2014.

Laurence, Stephen, and Eric Margolis. "Concepts and Cognitive Science." In *Concepts: Core Readings*, edited by Eric Margolis and Stephen Laurence, 3-81. Cambridge, MA: MIT Press, 1999.

Levine, Joseph. "Experience and Representation." In *Consciousness: New Perspectives*, edited by Quentin Smith and Aleksandar Jokic. Ashgate, 2003a.

—. "Knowing What It's Like." In *Privileged Access: Philosophical Accounts of Self-Knowledge*, edited by Brie Gertler, 45-53. Ashgate, 2003b.

—. "Materialism and qualia: The explanatory gap." *Pacific Philosophical Quarterly* 64 (October 1983): 354-61.

—. "On the Phenomenology of Thought." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 104-119. Oxford: Oxford University Press, 2011.

Lycan, William. "The Case for Phenomenal Externalism." *Philosophical Perspectives* 15 (2001): 17-35.

Margolis, Eric, and Stephen Laurence, . *Concepts: Core Readings*. Cambridge, MA: MIT Press, 1999.

—. "Concepts." *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2014/entries/concepts/>>.

Medin, Douglas, and Marguerite Schaffer. "Context Theory of Classification Learning." *Psychological Review*, 1978: 207-238.

Montague, Michelle. "Cognitive phenomenology and conscious thought." *Phenomenology and the Cognitive Sciences*, 2015: 1-15.

Pautz, Adam. "Sensory Awareness is not a Wide Physical Relation: an Empirical Argument against Externalist Intentionalism." *Nous* 40, no. 2 (2006): 205-240.

—. "The Real Trouble for Phenomenal Externalists: The Science of Taste, Smell, and Pain." forthcoming.

Peacocke, Christopher. *A Study of Concepts*. MA: MIT Press, 1992.

- Pitt, David. "Intentional psychologism." *Philosophical Studies*, 2009: 146 (1):117 - 138.
- . "Introspection, Phenomenality, and Content." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 141-173. Oxford: Oxford University Press, 2011.
- . "The phenomenology of cognition Or What is it like to think that P?" *Philosophy and Phenomenological Research* LXIX (2004): 1-36.
- Prinz, Jesse. *Furnishing the Mind: Concepts and their Perceptual Basis*. Cambridge, Massachusetts: MIT Press, 2002.
- Rosch, Eleanor, and Carolyn Mervis. "Family Resemblances: Studies in the Internal Structures of Categories." *Cognitive Psychology*, 1975: 573-605.
- Rosenthal, David M. "Sensory Qualities, Consciousness, and Perception." In *Consciousness and Mind*, by David M. Rosenthal, 175-226. Oxford: Oxford University Press, 2006.
- . "The Independence of Consciousness and Sensory Quality." *Philosophical Issues* 1 (1991): 15-36.
- Siewert, Charles. "Is Experience Transparent?" *Philosophical Studies* 117 (2004): 2004.
- . "Phenomenal Thought." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 236-267. Oxford: OUP, 2011.
- . *The Significance of Consciousness*. Princeton, NJ: Princeton University Press, 1998.
- Thompson, Brad. "The Spatial Content of Experience." *Philosophy and Phenomenological Research*, 2010.
- Tye, Michael. "Inverted earth, swampman, and representationalism." *Philosophical Perspectives*, 1998: 12 (S12):459-78.
- . "Phenomenal Externalism, Lolita, and the Planet Xenon." In *Collection on the Philosophy of Jaegwon Kim*, edited by Terence E. Horgan & David Sosa (eds.). forthcoming.