

Can Phenomenology Determine the Content of Thought?

Peter V. Forrest
Wolfson College
D.Phil Candidate in Philosophy

Abstract:

This thesis is about consciousness and representation. More specifically, the big picture issue in the background throughout is the *relationship* between consciousness (or “phenomenology”) and representation (or “intentionality”) in the life of the mind.

Phenomenology and intentionality are inarguably the two central topics in philosophy of mind of the last half-century. The question of phenomenology is, “how can there be something it feels like, from a subjective viewpoint, for a physical being to experience the world?” The question of intentionality is, “how can something physical, such as a brain state, be about, or represent, some other thing out in the world?”

Not too long ago, the majority opinion was that these two questions addressed two essentially independent domains. However, in recent years the views of many philosophers have swung dramatically in the opposite direction. An important theme of analytic philosophy of mind in the last decade or two has been the exploration of the groundbreaking idea that these two domains might be fundamentally linked in previously unrecognized ways. Perhaps phenomenal properties are reducible to certain kinds of intentional properties. Perhaps the mind’s non-derivative intentionality is grounded in phenomenology. Perhaps we should think of phenomenology and intentionality as “intertwined, all the way down to the ground” (Chalmers 2004, 32).

This thesis addresses one crucial question within this larger framework: whether, and how, thoughts are phenomenally conscious. Thoughts are an important test case for theories about the relationship between phenomenology and intentionality, because they have long been considered paradigmatic intentional states, in contrast to perceptual and sensory experiences, which are paradigmatic phenomenal states. While there is something

it is like, from the inside, for an individual to undergo a perceptual experience such as an olfactory experience of roasted coffee beans, by contrast entertaining a thought might seem to lack such a distinctive qualitative “feel”. The thought is clearly intentional: it involves carrying informational content about objects and properties in the world. But is there also something it is like for a subject to experience thinking itself? To answer this question in the affirmative is to accept the existence of a phenomenology of thought, so-called “cognitive phenomenology” (CP).

The literature on this topic so far has focused primarily on the question of whether CP exists. Here I will focus on the subtly different, and largely neglected, question of whether a kind of CP exists that is able to determine thought’s intentional content. Many proponents of CP seem to be motivated by the hope that it can, since they believe that in the case of other conscious states, the phenomenology accounts for the intentionality. However, in what follows I argue that this ambitious project is doomed to fail, because CP is not suited to determine the intentional content of thought.

Table of Contents

Chapter 1: The Cognitive Phenomenology Debate.....	5
Chapter 2: Are Thoughts Ever Experiences?.....	55
Chapter 3: Phenomenal Contrasts and Self-Knowledge.....	82
Chapter 4: Cognitive Phenomenology and the Fringe.....	135
Chapter 5: Why Phenomenology Cannot Determine the Content of Thought.....	177
Bibliography.....	224

Chapter 1

The Cognitive Phenomenology Debate

I. Introduction

This thesis is about consciousness and representation. More specifically, the big picture issue in the background throughout is the *relationship* between consciousness (or “phenomenology”) and representation (or “intentionality”) in the life of the mind. Phenomenology and intentionality are inarguably the two central topics in philosophy of mind of the last half-century. The question of phenomenology is, “how can there be something it feels like, from a subjective viewpoint, for a physical being to experience the world?” The question of intentionality is, “how can something physical, such as a brain state, be about, or represent, some other thing out in the world?”

Not too long ago, so the story goes, the majority opinion was that these two questions addressed two essentially independent domains. The legacy of this way of thinking is still present in many introductory classes and textbooks in philosophy of mind today, where the standard narrative that undergraduate students are taught involves separate characterizations of phenomenology and intentionality, followed by independent overviews of the literature on the various problems posed by each in turn.

However, in recent years the views of many philosophers have swung dramatically in the opposite direction. An important theme of analytic philosophy of mind in the last decade or two has been the exploration of the groundbreaking idea that these two domains

might be fundamentally linked in previously unrecognized ways. Perhaps phenomenal properties are reducible to certain kinds of intentional properties. Perhaps the mind's non-derivative intentionality is grounded in phenomenology. Perhaps we should think of phenomenology and intentionality as "intertwined, all the way down to the ground" (Chalmers 2004, 32).

Even if these ideas ultimately prove to be mistaken, it is clear they cannot be ignored. On the contrary, our attempts to understand phenomenal consciousness and intentional content can no longer be carried out in isolation from one another. At the very least, the future of philosophy of mind must grant a central place to the question of how phenomenology and intentionality relate to each other in a single, unified theory of mind.

This thesis addresses one crucial question within this larger framework: whether, and how, thoughts are phenomenally conscious. Thoughts are an important test case for theories about the relationship between phenomenology and intentionality, because they have long been considered paradigmatic intentional states, in contrast to perceptual and sensory experiences, which are paradigmatic phenomenal states. While there is something it is like, from the inside, for an individual to undergo a perceptual experience such as an olfactory experience of roasted coffee beans, by contrast entertaining a thought might seem to lack such a distinctive qualitative "feel". The thought is clearly intentional: it involves carrying informational content about objects and properties in the world. But is there also something it is like for a subject to experience thinking itself? To answer this question in the affirmative is to accept the existence of a phenomenology of thought, so-called "cognitive phenomenology" (CP).

The literature on this topic so far has focused primarily on the question of whether CP exists. Here I will focus on the subtly different, and largely neglected, question of whether a kind of CP exists that is able to determine thought's intentional content.¹ Many proponents of CP seem to be motivated by the hope that it can, since they believe that in the case of other conscious states, the phenomenology accounts for the intentionality. In this way CP is supposed to provide the key to a theory of the foundations of thought, by explaining how thought states possess intentionality in the first place. However, in what follows I argue that this ambitious project is doomed to fail, because CP is not suited to determine the intentional content of thought.

2. In order to move forward, we need to get a bit clearer on what cognitive phenomenology is supposed to be. The recent and ongoing debate in the philosophical literature over the existence of CP has succeeded impressively in polarizing philosophers of mind into two diametrically opposed, equally ardent factions. Supporters of CP insist that cognitive phenomenology is an everyday feature of their conscious thought life and is manifestly obvious to introspection. CP skeptics deny that they have ever been aware of any such thing. Both sides of the debate respond with incredulity to their opponents' claims.²

One might be forgiven for being surprised by this level of disagreement (well, not if one has spent any length of time around philosophers). For one might suppose that the

¹ An exception to this neglect is Pautz 2013.

² Examples of the former group include Horgan and Tienson (2002), Pitt (2004), Siewert (1998), and Strawson (1994). Examples of the latter include Carruthers and Veillet (2011), Lormand (1996), Prinz (2011), Robinson (2005), and Tye & Wright (2011). See the contributions in Bayne and Montague (eds.) 2011.

question of whether phenomenology of thought exists is trivial, since if the question is whether there is something it is like, from the subjective viewpoint, to think, the answer is obviously “yes”. When I think about my mother, I experience a visual image of her face in my mind’s eye, and when I think about what I should say next in this paragraph, I experience the imaginary “sound” of my inner voice. So if there is something it is like to see and hear, then similarly there can be something it is like to think.

But in these examples, what it is like to think, the visual and auditory imagery, is in an intuitive sense “borrowed” from the perceptual experiences of seeing my mother and hearing my own voice. The phenomenology of the imagery experiences seems to resemble the phenomenology of the perceptual experiences. So philosophers typically do not consider mental imagery to provide evidence for cognitive phenomenology. Rather, the question for them is whether there exists some further, distinctive phenomenology that thought possesses “all to itself.” In other words, the debate over cognitive phenomenology concerns CP in the following sense:

Cognitive Phenomenology (CP): a *sui generis* type of phenomenology that is “proprietary” to thought

This initial definition of cognitive phenomenology raises three important questions, which will structure our discussion for the remainder of the chapter. First, what is phenomenology? Second, what does it mean for phenomenology to be *proprietary* to thought, i.e. how must phenomenology be related to thought to qualify as CP? Third, what does it mean for phenomenology to be proprietary to *thought*, i.e. what understanding of thought or cognition is operative in the CP debate? In the following sections, we will address each of these questions in turn.

II. Phenomenology

For a mental state to have *phenomenology* or *phenomenal character* is for it to be conscious in a particular sense, wherein to call a state conscious does not necessarily imply that the subject is conscious *of* the state. Rather, for a mental entity (state, process, event, etc) to be phenomenally conscious simply consists of there being *something it is like* for the subject to be in that mental state, i.e. to undergo the experience. The phenomenal (qualitative) character of the state constitutes what it is like for the subject. Unless otherwise specified, by ‘conscious’ I will always mean phenomenally conscious.

I understand experiences to be all and only those mental entities (states, events, or processes) with phenomenology (instantiating phenomenal properties). So ‘experience’ is basically synonymous with ‘phenomenal state,’ and every experience is a conscious experience on my view. I take the term ‘qualia’ to be synonymous with ‘phenomenal properties,’ but due to the controversy surrounding this term, I will avoid using it. However, I am not hesitant about describing all phenomenal character as *qualitative* character, in the sense that phenomenal properties always constitute the unique, ineffable, subjective, conceptually irreducible nature of a given experience, which distinguishes it as the particular type of phenomenal state that it is. Being qualitative in my sense does *not* imply that phenomenal character is essentially *non-intentional*. One of the key marks of being qualitative is that our grasp of such properties is conceptually independent from, and not perspicuously reducible to, properties grasped by functional or structural concepts. In other words, that there be an “explanatory gap” between our understanding of the former and the latter sorts of properties (Levine 1983).

Can phenomenology be characterized in a more detailed and illuminating way? I propose we can go beyond this standard characterization by accepting the promising thesis that phenomenal character is *presentational* in nature. The idea 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 objects) with determinate qualities. I think this is the notion that Kant was after when he spoke of “intuition”. Intuitions for Kant, as I understand it, are what provide the understanding with the concrete, singular material to which to apply concepts. 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” or “appearance” nature of phenomenal character certainly seems accurate for perceptual experiences generally. Perhaps it is more difficult to evaluate his claims when applied to alleged cases of non-sensory-perceptual experience, though Levine explicitly intends them to generalize to every instance of phenomenal character.

One strong consideration in favor of believing that this characterization does generalize starts from the standard gloss of phenomenology as “what it’s-likeness.” For 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.³ This fact suggests a certain structure inherent to experience: there is (1) the experiencing subject, (2) an apparent two-place relation of awareness connecting subject and object, and (3) the property or properties that constitute what it is like for the subject, i.e. the “objects” of the subject’s awareness.⁴ To say that an

³ Cf. Crane 2014, 162-163.

⁴ Someone might complain that this characterization of phenomenology does not adequately differentiate phenomenal states from intentional states of mind generally, since arguably there are non-phenomenal

“object” is “presented” or “appears” to a subject in experience is just to say that the subject enjoys this direct phenomenal access—sometimes dubbed “acquaintance”—to one or more phenomenal properties. In virtue of this awareness the subject typically can direct her attention to these properties, though arguably, the minimal state or relation of phenomenal awareness is prior to the conscious, willful directing of attentional resources. Here is the Presentation thesis:

Presentation: all phenomenology is presentational in nature, consisting exclusively in the presentation of apparently distinct objects (i.e. particulars and properties) to the subject 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.

It is important to clarify what this claim about the nature of phenomenology does *not* amount to. For a subject to be presented with “objects” of experience in this way does not entail that we experience the object *as* an object, or apply a representation to the experience that means *object*—at least not in any robust sense of objecthood that, e.g., meets the standards of our intuitive theory of physics. Nor does the claim obviously imply that the objects of which we are phenomenally aware are mind-independent, objective, public, concrete, or external entities. And because this notion of the presentational nature of phenomenology 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

intentional states that are equally “for the subject,” i.e. have something like a subject-object structure. But I am not attempting to provide a set of necessary and sufficient conditions for phenomenal states. I would prefer not to describe non-phenomenal intentional states as “presentational” in the sense described above, since I intend to use this term only for the kind of immediate, acquaintance-like awareness of properties that constitute what it is qualitatively like from the first-person perspective, and in my book not every personal-level state of “awareness” of intentional content counts as presentational in this strong sense. But I recognize that to defend this distinction thoroughly would require a thorough analysis of “what-it’s-like” talk and/or “acquaintance” talk, which I am not prepared to give here. Since I am not attempting to give a definition of phenomenology, but rather just a characterization that will allow us to get an intuitive grip on the notion, these issues can be set to one side.

commitment to any particular account of the underlying metaphysical nature of this phenomenal awareness, for instance whether we should prefer a sense data or direct realist theory of experience rather than an adverbialist theory.

My view *is* in line with Christopher Frey's recent proposal that, "when we phenomenally appreciate a [feature of] an experience, we appreciate it as being both something other than ourselves and as standing in opposition to ourselves," and never as "being, being instantiated in, or being about the self qua experiential subject (or a state/mode thereof)" (2013, 76). One might question whether some purported types of phenomenology, e.g. in the experience of agency, necessarily involve experiencing their features as "other than ourselves." However, the qualifier "qua experiential subject" is crucial here, since for phenomenology to be presentational does not rule out the possibility of idealism, or that the phenomenological evidence may be neutral as to whether the properties of which we are directly aware in experience could exist independently of the subject. The idea, rather, is that the basic awareness of phenomenal character constitutive of every conscious experience involves a *distance*, noticeable on introspective reflection, between the experiencing subject and what is being experienced.⁵ This structure of awareness of determinate objects other than our experiencing selves is internal to every possible conscious experience.⁶

⁵ One might say this amounts to the claim that phenomenology always involves an apparent "act-object" structure. But I am wary of using terminology from the history of philosophy that comes with so much baggage. So I prefer instead to put things in my own terms, and let readers draw connections to past views and debates if they wish.

⁶ One might suppose that this view of phenomenology's basic structure commits us to denying that there can be phenomenal properties that are non-intentional "raw feels" or "mental paint." (Levine (2011, 111) says as much.) If we define intentionality as the mind's "directedness" towards an "object", then Presentation does entail that phenomenology brings with it intentionality. For if endorsing the existence of "raw feels" essentially involves denying this minimal sort of "directedness," then on my view the notion of "raw feels" would be unstable: the "feel" is just what it's like for the subject, and there can be no "feel" without its being

However, the Presentation thesis is independent from any suggestion that experience's built-in subjectivity, its "for-me-ness", is itself an element of the experience's overall phenomenal character. I find this view about phenomenal subjectivity, which is advocated by those who hold that phenomenal experience is pre-reflectively "self-representational," to be implausible.⁷ But more to the point, it does not follow simply from the observation that all phenomenology is "for the subject", and therefore involves a presentational, subject-object, structure. A subject certainly could be aware of herself, in some guise or other, as a phenomenal feature of her overall conscious experience. But such self-awareness clearly is not entailed by the bare presentational fact that, upon reflection, we can realize that all experience must involve a certain structure that presupposes a subject in the minimal sense of requiring a *subjective point-of-view*. Nor does this realization imply that subjectivity in this guise is an introspectable feature of what it is like to undergo the experience, in the very same way that e.g. phenomenal blueness and the smell of garlic are introspectable features. Indeed, the suggestion that the subject *qua experiencer* could also be the *object* of her own experience is hard to wrap one's head around, and may be just incoherent.

Someone might worry that the presentational account of phenomenal consciousness I have just described provides an overly narrow "perceptual model" for understanding all phenomenology. But the primary motivation for the view has come simply from our

subjectively "felt," which is just the subject being directed on the phenomenal feeling in the way we have been discussing. So it would be incoherent to claim that there could be a "feel" that is non-intentional in the sense of lacking directedness. But we do not need to assume the presence of intentionality in any more robust sense, e.g. that the experience involves a seems/is distinction, accuracy or truth conditions that genuinely could fail to be met, or that it is directed to things beyond itself. There might be a straightforward and compelling argument from the observation that phenomenal character is presentational to the conclusion that it is essentially intentional in a more robust sense. But I think the latter claim will require additional argumentation.

⁷ See the contributions in Kriegel and Williford 2006.

consideration of what follows from there being something it is like for a subject to undergo an experience. Since we cannot assume that what-it's-likeness is limited to perception without begging the question against the defender of CP, we also should not assume that the presentational structure within experience is limited to perception.⁸

III. Two CP Theses

1. We can turn now to the second question that arises from our initial description of cognitive phenomenology: what does it mean to claim that cognitive phenomenology is *proprietary* to thought, i.e. how must phenomenology be related to thought to qualify as CP?

I think there are two primary ways to answer this question, and this thesis is in many ways structured around the idea that distinguishing between these answers is of central importance to the CP debate. On the one hand, someone who adopts what we can call a *Modal Thesis* will interpret the claim that phenomenology is 'proprietary' to thought to imply that phenomenally conscious thoughts possess cognitive phenomenal properties that, as a matter of metaphysical necessity, are unique to thoughts. This in turn implies that the instantiation of such phenomenal properties in the mind of a subject is sufficient to determine the presence of an occurrent, conscious thought in this subject. Cognitive phenomenal properties are thus inseparable from thoughts: across all possible worlds, there is no cognitive phenomenology in a subject's mind without an accompanying cognitive state (though not necessarily vice versa). For instance, this inseparability would

⁸ Of course, the view does not claim that we are aware of phenomenal properties *in the very same way* that we are aware of external physical objects. I do not literally *see* the phenomenal qualities of my visual experiences, for instance.

not hold in the case of phenomenal red (the phenomenal property instantiated in normal subjects when they visually perceive a red object), since, plausibly, phenomenal red can be experienced by a perceiving subject without that subject thereby entertaining any thoughts at that time. So phenomenal red is plausibly not a cognitive phenomenal property in the Modal sense.

The Modal Thesis does not say anything about what it is like to have an experience with CP character, what it “feels” like, qualitatively, to think. For all we have said so far, perhaps it is like being tickled, or having indigestion. I just argued that phenomenal red is not a CP property, but nothing about the Modal Thesis *per se* explicitly rules out sensory-perceptual phenomenology, such as color phenomenal properties, from counting as CP. On the other hand, we may want to endorse a *Qualitative Thesis*, according to which the idea that CP is proprietary to thought implies that (at least some of) the phenomenology of (at least some) thoughts is *qualitatively different* from other familiar phenomenal kinds. The key idea of the Qualitative Thesis is simply that if we accept CP then our ontology must be slightly expanded: contrary to what many have supposed, phenomenology is not exhausted by sensory phenomenology (broadly understood to include the phenomenology of low-level perceptual experience and sensory imagery). Thought need not borrow its qualitative, what-it’s-like character from another mental domain, because it has its own.

However, the Qualitative Thesis does not entail that the instantiation of cognitive phenomenology is metaphysically sufficient to ensure the presence of conscious thought. This means that, at most, the connection between cognitive phenomenology and thought is one of nomic necessity, but it may be that the connection is even weaker. For instance, it may be that, as a matter of fact, in humans CP is only ever present in the mind when

accompanied by a thought, but that it would nevertheless be possible for the capable mad scientist to induce an experience of cognitive phenomenology in a subject without any corresponding thinking. Indeed, as far as the Qualitative Thesis is concerned, it may even be possible to experience cognitive phenomenal properties without thinking in everyday life (though such experiences would be the exception that proves the rule). All that the thesis commits one to is the claim that there are phenomenal properties that always or typically accompany a normal subject's experiences of thinking, and/or that play an important role in relation to our faculty of thought, *and* that are qualitatively different from any kind of sensory phenomenology.

Therefore, to move forward on the question of cognitive phenomenology, we need to acknowledge these two versions of the claim that CP exists, based on two ways of interpreting the notion that CP is proprietary:

Qualitative Thesis: there exist phenomenal properties that are qualitatively different from all sensory phenomenal properties, and which are (in some way) specially associated with thought

Modal Thesis: there exist phenomenal properties that are *necessarily* unique to, and inseparable from, thought: the instantiation of one of these phenomenal properties in a subject is sufficient (in every possible world) to determine the conscious representation of a thought content, i.e. the tokening of a conscious thought in the subject

Based on these two CP theses, we can distinguish between four positions that one might adopt on the issue of cognitive phenomenology:

	Qualitative Thesis	No Qualitative Thesis
Modal Thesis	Strong CP	Empiricism?
No Modal Thesis	Moderate CP	Anti-CP

First, one might deny both the Modal and the Qualitative Theses and advocate eliminativism with respect to all forms of cognitive phenomenology. I reserve the label *Phenomenal Conservative* for those that choose this position.⁹ Not only do Conservatives deny the existence of CP, they also typically restrict the class of phenomenal properties to the sensory realm. Second, one might accept the Modal Thesis but reject the Qualitative Thesis, in which case one's position is a form of empiricism rarely defended in contemporary philosophy: on this view, phenomenology is entirely sensory, but nevertheless having certain sensory experiences is sufficient to determine that one is thinking. Third, one might accept both the Qualitative and Modal Theses, in which case one accepts the existence of *Strong CP*, which is any phenomenal property the existence of which would render the conjunction of the Qualitative and Modal theses true. Finally, one might accept just the Qualitative Thesis, in which case one accepts only *Moderate CP*, which is any phenomenal property the existence of which makes the Qualitative Thesis true but does not make the Modal Thesis true. All those who accept the existence of CP in some form or other I call *Phenomenal Liberals*.

Most avowed Liberals seem to support Strong CP, though they often do not clearly delineate the two CP theses that make up this position. Both notions, Modal and Qualitative, are clearly in play in the literature on the topic, from philosophers on both sides.¹⁰ The distinction between Modal and Qualitative Theses is somewhat obscured by

⁹ I am aware that the label 'phenomenal conservatism' has recently been used to pick out a very different distinction in epistemology (Tucker 2013). However, since that issue is unrelated to the discussion here, and since "conservative" and "liberal" are easy, intuitive handles for the two sides of the CP debate, I have kept this terminology (Prinz and Siewert opt for the much more awkward "restrictivist" vs. "expansionist" and "exclusivist" vs. "inclusivist" labels, respectively).

¹⁰ For clear expressions of (something like) the Qualitative Thesis, see Bayne ms, Prinz 2011, Siewert 2011, and Smithies 2012a. For clear expressions of (something like) the Modal Thesis, see Horgan and Tienson 2002, Pitt 2004, and Tye and Wright 2011.

the fact that one of the most common Conservative responses to Liberal arguments is to argue that all of the apparent evidence adduced in support of CP can be accounted for entirely by the presence of some kind of sensory phenomenology. This in turn has led Liberals who accept Strong CP to focus on defending their position by arguing that their opponents' attempts to explain the data in terms of sensory phenomenology alone are inadequate. But this argumentative strategy, if successful, only *directly* supports Moderate CP.

In this thesis, I am interested primarily in refuting the Modal Thesis, the claim that there is a kind of phenomenology that metaphysically determines the representation of a thought content, i.e. the thinking of a thought. So I cannot agree with most Phenomenal Liberals who affirm the existence of Strong CP. A secondary aim of the thesis is to explore and defend a version of Moderate CP. So I also differ from Phenomenal Conservatives, who reject all forms of CP and often endorse the view that all phenomenology is sensory.

2. Drawing the distinction between Modal Thesis and Qualitative Thesis allows us to characterize cognitive phenomenology in terms of either Strong CP or Moderate CP. But note that doing so still leaves many questions regarding the nature of CP wide open. For instance, how detailed and various is cognitive phenomenology? One might think that there is a type of phenomenal property that is common to every experience of thinking or understanding—or at least is common to many different experiences of thinking across differences in intentional content or attitude type. One might even think that all cognitive

phenomenology is exhausted by this type of phenomenal character: the phenomenology of *thinking in general*.

Alternatively, one might recognize a distinct phenomenal type for each propositional attitude state type, which is common and unique to conscious tokens of that attitude type. On both this view and the previous view, what it's like to judge that my wife is beautiful is the same as what it's like to judge that Coventry is ugly, at least with respect to the cognitive phenomenology of the two thoughts. On these views CP is too course-grained to differentiate thoughts on the basis of differences in their intentional contents.

In contrast, we might adopt a position according to which cognitive phenomenology varies with differences in the intentional content of each cognitive state. Each distinct thought type, distinct in content, has qualitatively different cognitive phenomenology. My thought that oranges are delicious has different phenomenal character than does my thought that apples are delicious (and note this could be true whether we defend Moderate or Strong CP).

Assuming that CP is distinctive for different thought contents in the way just described, one may go on to wonder whether cognitive phenomenal properties are *intersubjectively* or *intrasubjectively* distinctive: for instance, must my thought that water is wet have the same cognitive phenomenal properties as your thought that water is wet, or is it enough that each of *my* thoughts about water being wet has the exact same cognitive phenomenology, irrespective of what other people's water thoughts are like? And must what it's like for me to think that water is wet remain constant over periods of time?

Our definitions of Moderate CP and Strong CP also leave open the question of whether cognitive phenomenal experience is *pure*, *impure*, or a combination of both. A

token CP experience is pure just in case its phenomenal character is exhausted by cognitive phenomenal properties; otherwise it is impure (Levine 2011). One might endorse Moderate or Strong CP, but think that all CP is necessarily impure, perhaps because one thinks that a subject cannot undergo an experience with CP character in the absence of sensory imagery.

Finally, characterizing CP in terms of the Modal Thesis leaves open the question of the exact metaphysical relationship between phenomenal character and thought content, or between experience and thought. Is the relationship one of metaphysical *dependence*, i.e. does cognitive phenomenology “ground” the representation of thought content (or vice versa)? Or should CP character be conceived of as *identical* to a kind of intentional content? Or perhaps the fact that a CP property is never instantiated in a subject’s mind in the absence of a corresponding thought is just a brute necessity, without further explanation. As I am using these terms, the mere claim that CP metaphysically “necessitates” or “determines” thought is silent as to which of these three options is correct. So the Modal Thesis, the claim that cognitive phenomenology determines thought, does not tell us *why* or *how* it does so. But we will return to this question in Chapter 5.

3. Sensory and Non-Sensory Experience

The Qualitative Thesis claims that there exists a type of phenomenology that is qualitatively distinct from all sensory phenomenology. This claim raises the obvious question: how should we understand the alleged difference between sensory and non-sensory phenomenology? Before moving on to consider thought in the next section, we need to sketch an answer to this question.

The sort of contrast we need to describe is phenomenal in the sense that it marks out two different types of phenomenology purely on the basis of a difference in what it is like for the subject, which holds between every property of the first type and every property of the second type. That is to say, there should be an uncontroversial, intuitive sense in which what it is like for a subject to experience a sensory phenomenal property is more similar to what it is like to experience every other sensory phenomenal property than it is what it is like to experience any non-sensory phenomenal property. Our model here is the intuitive differences we find between types of perceptual experience within and between the various sense modalities: just as any red experience is phenomenally more like every other red experience than like any green experience, and any color experience is phenomenally more like every other color experience than like any shape experience, and any visual experience is phenomenally more like every other visual experience than like any auditory experience, so too, any sensory experience should be phenomenally more like every other sensory experience than like any cognitive, or otherwise non-sensory, experience.

Still, how do we pick out which phenomenal properties are the sensory ones in the first place? Clearly, sensory phenomenology typically is associated with perceptual experiences and bodily sensations. So we might suggest that sensory phenomenal properties are just all and only those properties that characterize what it is like for a subject to be in some perceptual state (broadly construed to include bodily sensations). But I think it is important for us to keep the sensory/non-sensory *phenomenal* distinction independent from any perceptual/non-perceptual distinction. For some philosophers believe that perceptual experiences of so-called “higher-level” properties, such as kind

properties, involve a distinctive phenomenology, over and above the phenomenology of color, shape, motion, and other so-called “lower-level” properties represented in perception. As we will discuss in ensuing chapters, how to classify such higher-level perceptual phenomenology is currently an open and controversial question. Some philosophers will classify it as non-sensory, and of those some may consider it a type of cognitive phenomenology. Thus, in articulating a sensory/non-sensory distinction, we should not require from the outset that all perceptual experiences are sensory experiences, or that experiences have to be qualitatively distinct from “higher-level” perceptual experiences in order to count as qualitatively distinct from sensory experiences.

In any case it is clear that drawing the distinction between sensory and non-sensory experience along the cognitive/perceptual divide would not capture the common usage of the term ‘sensory’ by the various parties to the CP debate. Rather, when Phenomenal Liberals claim there are experiences with phenomenology that reaches beyond the sensory, and when Phenomenal Conservatives attempt to explain away the purported evidence for such experiences entirely in terms of sensory experiences, they each reach for the classic examples: color, shape, sound, pain, and so on. Thus it is pretty clear that both sides have in mind a characterization of sensory experience that aligns closely with what are called “lower-level,” “thin,” or “sparse” perceptual experiences. Such experience is limited to representing the sorts of properties that traditional empiricists believed were represented by perception. So in this thesis I will understand sensory experiences as all and only lower-level perceptual experiences, i.e. those perceptual experiences that exclusively represent lower-level properties. Sensory phenomenology is just what it is like to undergo a sensory experience.

What are lower-level perceptual experiences? I think the best way to proceed is to elucidate the relevant notion by listing examples. One could either attempt to give an exhaustive list of all the sorts of experiences that count as sensory in our sense, or else one could provide paradigmatic examples—e.g. visual experiences of color, space, motion; auditory experiences of pitch, timbre, and loudness; etc.— in the hopes that we all have an intuitive sense of those experiences that bear the requisite family resemblance to the paradigms. Whichever way we proceed, it is important that the phenomenology of so-called sensory imagery qualifies as sensory phenomenology, since no one writing about CP thinks that such imagery, when accompanying cognitive states, should count as genuine CP experience. But this qualification is unproblematic, since intuitively so-called sensory imagery *is* relevantly like low-level perceptual experiences, which is why it is so natural to describe imagery as “copies” of the sensory phenomenology of perceptual experience.

In addition to the appeal to examples however, it would be nice if we were able to describe sensory experience by citing one or more distinctive characteristics that sets all lower-level perceptual experiences (thus, all sensory experiences) apart from other types of perceptual experience. But this task proves difficult. A few obvious suggestions along these lines can be ruled out straight away. For instance, philosophers sometimes distinguish perceptual experiences from “raw,” or “mere,” sensations. The former serve to represent a subject’s immediate environment while the latter (allegedly) do not, but merely consist of non-intentional qualitative “feels”. But this will not work to capture our distinction, because many sensory experiences in our intended sense are unquestionably and essentially representational: think for instance of a simple auditory experience that represents a sound of a certain pitch and volume originating from somewhere over your

left shoulder. So being sensory in our sense is not the same as being a raw sensation. Recently a number of philosophers have been occupied with articulating a related distinction, viz. that between merely sensory experience on the one hand—understood as at most representing subjective, mind-dependent properties—and, on the other hand, full-blown perceptual experience—understood as representing an objective, mind-independent world.¹¹ While this distinction is an important one, again it will not suit our purposes here, and for the same reason: many sensory experiences in our sense, e.g. the visual experience of a blue triangle painted on the wall before me, are clearly presented as objective, mind-independent features of our external environment.

A more promising idea of how to characterize sensory experience comes from Jesse Prinz. In his introductory remarks on the cognitive phenomenology debate, Prinz suggests that we should understand the sensory properties of experience as those properties represented in experience that constitute “some aspect of appearance,” and non-sensory properties as those that “transcend appearance.” He speaks in similar terms of sensory experiences representing “superficial features” of the world, and non-sensory properties as features that “go beneath the surface” and are “indistinguishable by the senses” (2011, 175-176).

There is an intuitive attraction to characterizing sensory experiences as “appearances” in this way. But Prinz’s suggestion nicely illustrates the difficulty inherent in trying to articulate a characteristic that sets all and only sensory experiences apart. For in order for his characterization to be helpful, he owes us a more specific account of which

¹¹ For variations on this theme, see Siegel 2006b, Burge 2010, Matthen 2010, Farkas 2013, and Masrour 2013.

properties are presented to us in appearances.¹² For a property to appear to us, it does not need to be literally *on the surface* of the things we perceive, since it is possible that a sense organ could pick up on properties within a physical body (the ability of some snakes to sense the thermal radiation of other animals may provide an example). Nor should we say that a property represented in perception counts as appearing to us only if we can detect its presence in our environment without needing to draw any inferences, or only if it is somehow immediately “given” to our senses at the level of our sensory transducers. For many of the low-level sensory properties that we want to count as appearing to us—e.g. visual experience of depth—are known to require sub-personal inferential processes, including feedback from points higher up in the chain of perceptual processing, before they ever show up in conscious experience (cf. Fodor 2007, Masrour 2011).

Intuitively, the properties that appear to us, in Prinz’s sense of representing “superficial features” of the world, are just those that we perceive—that we can see, hear, feel, and so on. But if this is taken to pick out all properties represented in perception we have a problem, since arguably perception represents many properties that do not appear to us in any intuitive sense. So we might say instead that the appearance properties are those properties *consciously* represented in perception. But now, this way of specifying appearance properties, if combined with the assumption that only low-level sensory properties are included in appearance, begs the question against the “thick” theorist who claims that we perceptually experience such properties as being an object, being a cause,

¹² Earlier in this chapter I proposed we follow Levine in understanding *all* phenomenology as consisting in being “appeared to,” and explained this notion in terms of the Presentation thesis, which I motivated simply by appeal to the idea that all phenomenology characterizes what it’s (immediately, qualitatively) like *for a subject*. So whatever Prinz means by “appearances” here, it better be more demanding than mere Presentation, if it is to help us distinguish sensory from non-sensory phenomenology.

and being of a certain kind. Prinz has not provided an argument for thinking that such thick properties cannot “appear” to us as well. Ultimately he may turn out to be right on this point, but we want an initial characterization of sensory experience that is neutral with respect to the debate between Prinz and the thick theorist.¹³

Nevertheless, there is something attractive about the idea that sensory properties somehow are related to how things perceptually appear to us. I think that, while we cannot claim that all appearance properties are sensory, still the sensory does *constrain* perceptual appearance: *all perceptual appearance properties are either sensory properties or else their representation in perceptual experience involves a sensory basis*. Let me elaborate.

A key characteristic that separates sensory experiences from any other kind of perceptual experience, I want to suggest, is that the phenomenal properties of sensory experiences are the only perceptual phenomenal properties that are *phenomenally basic*. Intuitively, one could have a perceptual experience of thin properties, e.g. of colors and shapes, without experiencing any thick properties, e.g. of kinds; but one could not experience the kind properties without also experiencing some thin properties at the same time. So the phenomenal character of the experiences of the thin properties is more basic than the phenomenal character of the experiences of the kind properties. A phenomenal property (type) is phenomenally basic if and only if it partly or wholly constitutes the phenomenal character of a phenomenally basic experience. An experience (type) is

¹³ Furthermore, since being consciously represented in perception is equivalent to being represented in perceptual experience, Prinz’s empiricist claim that only appearance properties are represented in perceptual experience would carry more than a whiff of the tautological, when in fact this claim is understood to be substantive and controversial. So we better adopt a different characterization of what it means to be an appearance property.

phenomenally basic if and only if having an experience of that type does not necessitate having any other distinct (type of) experience.¹⁴

One might object to the claim that sensory phenomenal properties are phenomenally basic on the grounds that experiencing such properties often requires experiencing *other* sensory phenomenal properties, e.g. one cannot experience color phenomenology without also experiencing some spatial phenomenology.¹⁵ However, on the above definition, both color and spatial phenomenal properties are basic since they jointly constitute what it's like to have phenomenally basic *experiences* of color and space. These experiences, in turn, count as basic because there are no other experiences the having of which is required in order to have them: e.g. one does not need to have a color experience in order to have a color-and-space experience, for the simple reason that there is no such thing as *just* a color experience, devoid of spatial phenomenology.¹⁶ By contrast, having a visual experience of a thick property, such as being-a-llama, involves color and shape phenomenology, and thereby necessitates also having the simpler experience that consists of just the color and shape phenomenology. One could not have the llama experience without also having the color-and-shape experience, but, intuitively, one *could* have the color-and-shape experience without experiencing the llama, or indeed any other

¹⁴ Cf. Masrour's discussion of "phenomenological fundamentality" (2011, 22-23); also, Levine's description of what he calls the Kantian "sensory manifold" (2011).

¹⁵ Masrour (2011) raises this objection, but then in a footnote points out a similar way to respond to it.

¹⁶ This way of putting things does have the consequence (which I take to be completely innocuous) that there does not exist a distinct experience type for every distinct phenomenal property type: e.g. strictly speaking, there cannot be an experience whose phenomenal character entirely consists of a determinate shade of phenomenal red, since any such "experience", i.e. any instantiation of this red phenomenal property, is necessarily accompanied by the instantiation of some spatial phenomenal properties with which the red is bound up. So there is no such thing as *just* the experience of the red. Of course, we can still speak loosely of the red experience by itself, since we can distinguish the red phenomenology from the spatial phenomenology in thought.

higher-level kind properties. So the merely sensory experience of color and shape is phenomenally basic.

There are other higher-level properties, besides kind properties, that some philosophers believe we experience in perception. A moment's reflection suggests that these properties too—objecthood, causal relations, action affordances—are perceptually experienced, if at all, only along with low-level sensory properties like color and shape. I cannot see an object without seeing it as having certain dimensions, for example. So these examples also support the view that sensory experiences are the only perceptual experiences that are phenomenally basic. Indeed, if we did find a “higher-level” property that is represented in phenomenally basic perceptual experiences, then my bet is we would be inclined to classify the phenomenal character of such experiences as sensory anyway—along with sounds, shapes, colors, and the rest—since our intuitions about which properties bear a family resemblance to the sensory paradigms arguably track the facts about phenomenal basicness.

Let us sum up. We need a notion of sensory experience in order to have a sense of what cognitive phenomenology is supposed to look like according to the Qualitative Thesis. I have suggested that sensory phenomenology—as the phenomenal type with which CP is to be contrasted—corresponds to the phenomenology of what philosophers of perception call low-level or thin perceptual experience. At least, this seems to be the sense of ‘sensory’ that concerns Phenomenal Liberals and Conservatives when they argue over whether certain experiences can be entirely accounted for in terms of sensory phenomenology or not. And I also have argued that one characteristic of sensory experience that distinguishes it from non-sensory perceptual experience is that it is phenomenally basic.

One might worry that this claim unfairly limits what sort of phenomenology can count as CP before the debate has begun, since I have associated phenomenal basicness with sensory phenomenology, and the Qualitative Thesis, at least, stipulates that CP is non-sensory. So it may look like I have committed myself to the view that CP, if it exists, is not phenomenally basic. But I have been careful not to claim that only sensory phenomenology is phenomenally basic; rather, I think all and only sensory phenomenology is both phenomenally basic and *perceptual*. So there is space for the existence of *cognitive* phenomenology that is also basic in this way, i.e. phenomenology the experience of which does not necessitate any further, distinct experiences. It is true that this detail makes the characterization of sensory experience dependent on our account of the thought/perception distinction, since in order to determine if a given experience is sensory or not, one would need to know more than whether or not it is phenomenally basic; one would also need to know if it is perceptual or not.

But I do not see the phenomenal basicness constraint as providing a definition of sensory experience, nor as providing a fully sufficient guide to determining which phenomenal properties are sensory, if one doesn't already have a pretty good idea. The category of sensory experience is determined by our intuitive grasp of which experiences are relevantly like paradigmatic "low-level" perceptual experiences such as visual experiences of color and space, bodily sensations of temperature, pressure, and texture, auditory experiences of pitch, timbre, and loudness, and so on. In order to count as an instance of CP, an experience would have to be qualitatively different from all these sorts of experiences, whether or not the experience exhibited phenomenal basicness. The

phenomenal basicness criterion is only then introduced to explain what distinguishes all sensory phenomenal properties from other potential types perceptual phenomenology.

IV. Thought

1. What is Thought?

We have settled on our two formulations of the positive thesis of cognitive phenomenology, the theses defended by Phenomenal Liberals and rejected by Phenomenal Conservatives. The two formulations capture two different ways of understanding the Liberal position, depending on different ways we might interpret the claim that CP is *sui generis* or proprietary, namely a claim about its qualitative distinctness (Qualitative Thesis), or a claim about its modal tie to thoughts (Modal Thesis).

Both these theses raise the question of how we are to understand the notion of *thought* in the context of the cognitive phenomenology debate. In this section I will provide a characterization of the nature of thought, which should suffice for us going forward in the thesis. Along the way, I will explain my use of several related terms and how the mental categories they pick out—*cognition*, *perception*, *concepts*—relate to thought as I understand it.

What is thought? I want to begin by highlighting three guidelines we should respect in answering this question. *First*, we must not characterize thought in such a way as to unnecessarily bias or trivialize the CP debate from the start. *Second*, our description of the

nature of thought should be reasonably consonant with the core of what we already believe about thought, both in terms of our “folk” conception and in terms of what philosophers party to this debate seem to have in mind when they speak about “thought”. *Third*, our description of thought should clearly contrast it with perception.

The first two guidelines are obviously well motivated. The third guideline is motivated by the second, for in order to respect the views of thought held by philosophers who have written about cognitive phenomenology, we must account for the fact that the descriptor “cognitive” in “cognitive phenomenology,” is very often intended to be contrasted with “perceptual.” This tendency is hardly surprising, since philosophers and psychologists often distinguish between thought and cognition, on the one hand, and perception (and/or sensation), on the other. Furthermore, within the philosophy of mind, perceptual experience (understood broadly to include bodily sensations and the like) is the only paradigmatic, uncontroversial general category of phenomenal state, the type of phenomenal state that everyone tends to cite when introducing the notion of phenomenology to someone unfamiliar with the idea. Thus, it is natural to see the claim that phenomenology of thought exists as, in the first instance, a claim that there is phenomenology beyond the perceptual. And this in turn naturally leads to the question of how thought differs from perception.

With our three guidelines in place, I turn to my account of thought. The most obvious thing to say about thoughts is that they are paradigmatic *intentional* states of the mind: they are about, or directed at, things, typically things other than the states themselves. They have an intentional or representational *content* or subject matter, which

is what they are about, the information they convey.¹⁷ I adopt the common conception of all mental contents—including the contents of our conscious experiences—as *satisfaction conditions* or *conditions on extension* (Chalmers 2004; Siegel 2015). In particular, propositional thoughts—and propositional attitude states such as beliefs—have accuracy or truth conditions: their contents are just the conditions under which they are accurate or true.¹⁸ However, thoughts are not the only intentional mental states, and are probably not even the most numerous. Among many others, *perceptual* states, including perceptual experiences, are also undeniably intentional or representational. So we need to say more to describe what is importantly distinctive about thought that sets it apart from the rest of the representational mind.

To do so, I first want to provide a handful of paradigm examples of thoughts, to ensure that we all start more or less on the same page regarding the sorts of phenomena in question. Then I will delve deeper by outlining several general characteristics that all thoughts have in common.

In his summary article on cognitive phenomenology, Declan Smithies (2013a) has provided a list of everyday “cognitive activities,” what I am simply calling “thinking thoughts,” which I reproduce in part here:

¹⁷ This way of putting things blurs a useful distinction that is sometimes made between intentional objects (what the representation or intentional state is about) and intentional contents (how the object is represented as being, the particular way that it is characterized, by the intentional state), e.g. in Crane 2014. But this distinction will not matter for our purposes here. Also, I typically will use ‘intentional’ and ‘representational’ more or less interchangeably.

¹⁸ This conception as I adopt it is neutral as to whether these contents are (complexes of) worldly entities, i.e. the properties, objects, and states-of-affairs to which our words and mental states refer (Russellian contents); or instead should be thought of as modes of presentation of those entities, (Fregean contents); or just as sets of possible worlds; or whether our mental states should be thought of as having multiple kinds of content. I take it each of these views about the nature of contents is consistent with also thinking of them as being, or being given by, conditions on extension. The contents of other, sub-propositional intentional states are not truth or accuracy conditions, but just conditions for determining the objects or properties a state refers to, or is about.

- Considering a hypothesis
- Judging that a hypothesis is true
- Recalling a fact learned in the past
- Recognizing that the conclusion of an argument follows from its premises
- Inferring the conclusion of an argument from its premises
- Drifting aimlessly in thought
- Calculating the solution to a problem
- Deliberating about what to do
- Understanding a sentence
- Having a suspicion or a hunch (745)

In a similar vein, Elijah Chudnoff lists examples of what he calls “intellectual experiences,” which, I take it, count as intellectual in that they are instances of thinking: “intuiting that circles are symmetrical about their diameters, grasping a proof of the Pythagorean theorem, deciding to bike rather than walk to work, understanding what some passage is about” (2013, 182-183). Finally, Tim Bayne points out that,

With the possible exception of moments in which one is immersed in physical exercise, musical performance, or some form of meditation, the stream of consciousness is routinely punctuated by episodes of conscious thought. We deliberate about what to have for lunch, we remember forgotten intentions, we consider how best to begin a letter or end a lecture, and we puzzle over the meaning of a friend’s remark and the implications of a newspaper headline. (2011, 1)

I take it that considering these examples to be paradigms of thought is in line with our second guideline above: they are commonsensical examples of types of thought that we all enjoy every day, and they are offered by these authors as uncontroversial cases that those on both sides of the CP debate will be happy to accept, whatever else they disagree about.¹⁹

Painting in very broad strokes, we can discern considerable overlap in the examples: we might group them roughly into categories containing (a) instances of decision/judgment, (b) instances of remembering/recall, (c) instances of

¹⁹ Note the examples from Smithies and Bayne come from articles that are intended to be impartial, introductory summaries of the CP debate.

deliberating/puzzling/considering, and (d) instances of recognizing/infering/intuiting/grasping/understanding. Based on these examples, it is plausible to think the central type of thought is *judgment*, what is sometimes called an occurrent belief, which involves the taking up of an attitude of endorsement towards a propositional content. Judgments account for most thoughts in category (a). Arguably the examples in category (c), which involve extended activities of thinking rather than punctate events, can be understood largely as strings of such judgments. At least some examples from category (b), namely instances involving the recall of whole propositions from memory (e.g. suddenly remembering that you need to feed your cat), will be at least partly constituted by judgments, as will instances from category (d) of suddenly grasping the answer or meaning of a problem or message, provided the answer or meaning is propositional in form, and that grasping it involves tokening all of it in conscious awareness.

However, it is important to note that—even limiting ourselves to examples from the above lists—not every example of a thought with propositional content involves taking up an attitude of *affirmation* to the relevant proposition, e.g. instances of deliberating and considering can involve entertaining propositions without endorsing them as true. Nor is it obvious that every one of these examples involves a thought with propositional content. For example, one might contend that what we remember, recognize, infer, or understand is sometimes just an individual, rather than a proposition. We should leave these options open, keeping our conception of thought as broad and inclusive as possible.

I want to turn now to consider three general characteristics distinctive of all thoughts. In doing so, I will be describing thought (singular) *as a faculty*, rather than

speaking of thoughts (plural) as mental *particulars* (whether events, states, or processes), or of the *activity* of thinking. It is natural to conceive of thought as a unified faculty or capacity when characterizing what is distinctive about all thoughts, particularly in contrasting thought to perception and any other mental capacities that are functionally delineated by the cognitive sciences. The examples listed above, by contrast, are either instances of the activity of thinking, or are individual thought tokens, i.e. mental particulars. But the relationship between thought (faculty), thoughts (particulars), and thinking (activity) can be described straightforwardly, as follows. A subject exercises her faculty of thought just in case she engages in the activity of thinking. An episode of thinking, in turn, just amounts to the subject tokening one or more individual thoughts, simultaneously or in succession. (We also might want to insist that to count as a single episode of thinking the thoughts must be related in some non-arbitrary way, either by deductive or probabilistic inference or by association.) Throughout the thesis I will move freely between speaking of thought in these three ways.

The three general characteristics of thought are not so much individual characteristics as three families of similar features that thought exhibits. The first characteristic is that, in contrast to perception, thought's capacity to represent things is not limited in principle by facts about a subject's relations to her immediate, current environment. Call this characteristic thought's stimulus or ***Environmental Independence***.²⁰ The fact that my ability to think about an object or property is not dependent on any direct causal interaction between that thing and me means that what I am thinking about does not need to be in my nearby environment—it may not even exist.

²⁰ See Bayne 2013, 4-8. The discussion below is indebted in several places to points made by Bayne in his 2013, chapter 1.

Nor does the subject matter of my thoughts need to consist of those few types of objects and properties to which my sensory organs are sensitive and that my perceptual systems are equipped to represent. These facts have the potential to open up the horizons of the thinkable world, dramatically expanding it from the comparatively narrow world that perception discloses to the perceiver.

The flip side of this environmental independence is that the faculty of thought is not equipped to access information about the subject's environment directly, but can only make the subject aware of it indirectly, either via perception, or via inferences based at least in part on information stored in memory (which, presumably, ultimately depends on the past activity of perception). A slightly different way to put the overarching point is that the function of the perceptual modalities means they are "obligated to the immediate environment." When, for instance, a subject has a visual hallucination, when she sees what is in fact not there in front of her, we say vision has malfunctioned (Stokes 2013). Thought is not obligated to the environment in this way: we can think about an object being in front of us when it isn't without there being anything wrong with this thought, and even if we falsely *judge* that there is something in front of us when there isn't, it does not follow that the fault lies with our thought faculty.

The second main characteristic of thought is what we can call its ***Integrative Nature***: when producing new thoughts, our faculty of thought has the capacity, at least in principle, to draw on (the contents of) any other thought or any of the subject's beliefs and concepts, thus knitting together all the subject's thoughts and stored personal-level knowledge into an integrated whole. There are two primary strands to this characteristic. First, thought produces individual representations that are, in Stephen Stich's (1978)

phrase, “inferentially promiscuous,” meaning their contents stand in a great many (valid) inferential relations to one another and to the subject’s current belief contents, relations to which the faculty of thought is potentially sensitive, such that any thought could in principle lead to the subject’s inferring any other, and any thought or belief could in principle be seen to confirm or disconfirm any other. For instance, the implications of my thought that peaches are delicious are not necessarily isolated from my other (apparently quite disparate) beliefs and their implications, e.g. not only my beliefs about fruit, but also about the global economy, cognitive science, or God. Nothing about the powers of thought *per se* prevents me from recognizing semantic relations between this one thought and my other commitments quite generally, only provided I can make the connection. So thoughts are integrated with the rest of what the subject thinks and knows thanks to the availability of the contents of the former to be inferred from, and be the basis of inferences to, the contents of the latter.

Second, thought is integrative in virtue of the semantically evaluable *constituents* of individual thought representations relating in general and systematic ways, such that one’s ability to think certain thoughts guarantees the ability to think a great many other thoughts using the same constituents. If I can think that peaches are delicious, I should also be able to think that strawberries are delicious, or that peaches are expensive (provided I have the concepts of a strawberry and of expensiveness), since my use of the thought constituents meaning *peaches* and *delicious* ensures that I can entertain these same contents in the context of other thoughts. Thus, the ability to think about a particular thing (object, property, state-of-affairs) can be exercised quite generally, rather than only in the context of certain thoughts but not others. Of course, this point basically recapitulates Evans’s

famous “Generality Constraint” on concept possession, as well as the idea of “compositionality phenomena” famously discussed by Fodor: the compositionality, productivity, and systematicity of thought.²¹ But the main point for our purposes is simply that thought is integrated, not only in that the contents of individual thoughts stand in inferential relations to each other, but also in that the elements of thought contents can be combined and recombined to form a great many distinct thoughts. Furthermore, this integrative character enables a kind of unboundedness to thinking, accounting for our seemingly endless ability both to create novel thoughts from familiar constituents, and to draw novel connections between what otherwise seem to be unrelated ideas.

The final general characteristic of thought consists in a cluster of features that have to do with what we might call the *Metaphysical Profile* of the mental particulars that the faculty of thought generates. As we have mentioned already, the exercise of thought involves *intentional* (or representational) states that are about or directed at things, typically things other than themselves. Further, thoughts are *personal-level*, rather than sub-personal, states, meaning they are had by the whole person, or subject. A related

²¹ Evans (1982) describes the Generality Constraint as follows: “We cannot avoid thinking of a thought about an individual object *x*, to the effect that it is *F*, as the exercise of two separable capacities; one being the capacity to think of *x*, which could be equally exercised in thoughts about *x* to the effect that it is *G* or *H*; and the other being a conception of what it is to be *F*, which could be equally exercised in thoughts about other individuals, to the effect that they are *F*” (75); thus, “if a subject can be credited with the thought that *a* is *F*, then he must have the conceptual resources for entertaining the thought that *a* is *G*, for every property of being *G* of which he has a conception” (104). For the faculty of thought to be *productive* is for it to be able, in principle, to produce an infinite number of semantically distinct thoughts (presumably from finite representational means, i.e. the human mind can only contain a finite number of primitive representations, a finite “vocabulary” so to speak). For it to be *systematic* is for a subject’s ability to entertain certain thoughts to be systematically linked to her ability to entertain other thoughts. And for it to be *compositional* is for the content (i.e. meaning) of an individual thought to be determined by the contents of its constituent parts plus the syntactic rules for combining these constituents, such that the representations that serve as constituents make approximately the same semantic contribution to each thought of which they are a part (e.g. the representation that means *peaches* in my thought that peaches are delicious also must mean *peaches* in other thought contexts in which it finds itself, e.g. my thought that peaches are expensive). For further discussion see Fodor 1975; Fodor 1987, 135-154; Fodor 1997; and esp. Fodor and Pylyshyn 1988, 33-46.

feature, which may or may not amount to the same thing, is that all thoughts have the potential to be *conscious*. Obviously the question of whether thoughts can be phenomenally conscious is the principal point of controversy in this discussion, so the sense of ‘conscious’ here must not presuppose phenomenology, but only that thoughts should be the sorts of mental entities of which the thinker could be consciously *aware*.²² Finally, thoughts are *occurrent* entities. They happen or occur for the duration that they exist, in contrast to merely dispositional states of the mind, whose existence depends only on the mind’s propensity to produce certain behavioral or cognitive effects, given certain hypothetical conditions. It follows that, strictly speaking, most of our beliefs and other dispositional propositional attitude states are not instances of thought. Only so-called “occurrent beliefs” (desires, etc.) count as individual thoughts. But on my way of understanding them, these “beliefs” are actually no different than judgments, so it is easier to label them as such, saving the term ‘belief’ for a particular type of dispositional mental state. Beliefs, properly so-called, are stored in memory (or supervene on information stored in memory), but their propositional content, and the subject’s attitude towards it, can be expressed in thought by judgments, realizations, recollections, and the like.

Hopefully, the above characteristics of thought are sufficiently commonsensical and familiar to the reader as to seem fairly mundane and uncontroversial. I am not trying to reinvent the wheel here. I am also not offering a set of necessary and sufficient conditions for thought, but only a suitably interesting characterization of the subject matter that

²² We might understand this potential for conscious awareness in different ways, e.g. as referring to something like access consciousness in Block’s original (dispositional) sense (1995), or else in the weaker sense of merely being the type of representational state to which the subject is *able* to have “cognitive access,” in that the possibility is not ruled out by intrinsic facts about the representational type or by the cognitive structure of the subject’s mind, even if the thought is not presently “poised” to be accessed and used by the subject.

should be acceptable to all parties and can help to focus and clarify the CP debate. The everyday examples of thoughts offered at the beginning of this section are as important to this portrait as are the three characteristics, and should help guide the discussion when our three characteristics appear to underdetermine whether or not a given mental phenomenon should count as an instance of thought. To sum up, thought is that mental faculty whose capacity for representation (1) is Environment Independent, (2) is suitably Integrative, and (3) generates intentional, personal-level, potentially conscious, occurrent mental particulars, as found in (4) those examples listed in (a)-(d) above.

2. Thought, Cognition, & Perception

We still need to meet the third guideline I gave above for a successful account of thought: explaining how thought is meant to differ from perception. The notion of thought I proposed was meant to align with our folk conception. If we adopt a similarly commonsense notion of perceiving—seeing, hearing, feeling, and so forth—it is straightforward to see that the two differ. Most noticeably, perception, unlike thought, is not environment independent. It is “obligated to the immediate environment”—whether it counts as functioning correctly or not depends on how things are in the subject’s physical surroundings—and its capacity to represent depends on interactions between the sensory organs and proprietary types of proximal stimuli.

Further, perception arguably is not integrative to the degree that thought is. For one thing, there is the well-known fact that perceptual illusions, such as the Muller-Lyer and Ponzo visual illusions, persist even when they contradict what one consciously believes about what one is perceiving, a fact which has been taken as evidence that

perceptual states are in some sense isolated or “encapsulated” from the rest of what the subject believes (Fodor 1983).

Indeed, even if we remain neutral regarding any especially strong claims about the encapsulation of perceptual faculties, it is plausible that perceptual representation is not as inferentially integrated as is thought. It *may* be that some of my thoughts and beliefs could influence my visual faculty in such a way as to make me literally see the cat in my back garden as a tiger. But it is pretty unlikely that my beliefs about, say, the stock market, could directly affect how I perceive various objects and properties in my back garden. For even if we allow, for the sake of argument, that I can see the cat in my backyard as a tiger, and thereby perceptually represent tigerhood, it is highly doubtful that I could similarly perceive just *any* property that I can think about, e.g. that I could quite literally see my cat (or anything else in my visual field) as good-news-for-the-NASDAQ, and thereby perceptually represent this property. This limitation, however we spell it out in detail, puts a constraint on the sorts of inferential relations that can hold directly between perceptual states and the subject’s thoughts and beliefs, the kinds of direct influence that the one can have on the other and vice versa.²³

So even if there could be a tortuous inferential connection drawn between my current perceptual states and some highly unrelated set of my beliefs, such as my beliefs about the stock market, the faculty of perception itself could not draw the connection—it

²³ A related point is that the Integrative Nature of the semantically evaluable *constituents* of thoughts—which I discussed in relation to Evans’s Generality Constraint and Fodor’s “compositionality phenomena” in the last section—is not likely to apply analogously to the constituents of perceptual representations: just because I can perceptually experience the peach before me as sweet, it does not follow that I can perceive it as expensive, nor that I can perceive the NASDAQ as sweet. Likewise, of course, though I can taste the sweetness of the peach and see the redness of the strawberry, I cannot see the sweetness of the peach or taste the redness of the strawberry. These points follow from the domain specificity of the various perceptual systems that correspond to the different sense modalities.

would need thought to do that. For instance, I might be more likely to see the cat in my garden as a tiger if I believe I am in India, or that a tiger has recently escaped the local zoo; but in these cases, it is not likely to be either of these beliefs that directly affects what I perceive. Rather we require a judgment, inferred from these (and other) beliefs, to the effect that there could be a tiger loose in my back garden and thus present in my visual field. The faculty of thought (or unconscious cognition) would have to make the crucial inference to this additional thought, allowing what I see to be integrated with what I believe.

The upshot of these considerations is that a perceptual representation is not fully integrated with the rest of a subject's personal-level representational states to the degree a thought is, or, at the very least, *depends upon* the faculty of thought in order to be so integrated. In other words, to the degree that perception is connected up to the subject's "web of belief," it is so only from the periphery. Thus, the common conception of perception differs from our description of the faculty of thought at least with respect to both of our first two main characteristics of thought.

However, often in the psychological and philosophical literature, perception is primarily contrasted, not with thought, but with *cognition*. This is natural, for while 'thought' is a term expressing an everyday concept, the less common term 'cognition' suggests a more technical meaning. While earlier I attempted to describe thought in more-or-less commonsense functional terms, focusing on the nature of cognition invites us to consider in greater detail *how* thought does what it does, how its functions are carried out by particular "systems" and "mechanisms" (though still described at a fairly high level of abstraction above the neural level). Thus, the notion of cognition is most at home in

cognitive science, and the debate over where (or if) we should mark a sharp distinction between cognition and perception is naturally understood as a question to be answered not at the level of folk understanding or armchair reflection, but at the level of detailed, experimentally informed theory concerning how the human mind is actually organized.²⁴

But this leaves us with the question of how best to understand the relationship between thought and cognition. In the cognitive phenomenology literature, the two terms often seem to be treated as synonymous, as evidenced by the fact we speak of the alleged target phenomena both as “phenomenology of thought” and as “cognitive phenomenology”. But I think there are probably good reasons to distinguish the two. As I just suggested, *thought* is, first and foremost, a folk concept and *cognition* a theoretical, scientific one. As such, the domain of the latter arguably outstrips that of the former: every episode of thinking is (or strongly supervenes on) an exercise of one or more cognitive systems, but not every exercise of cognition involves thought in the narrow, commonsense way I have outlined.

For one thing, as I pointed out, many unquestionably cognitive states, including beliefs and many other propositional attitude states, are *dispositional* rather than occurrent, so do not count as thoughts proper. For another thing, many cognitive states and processes are *sub-personal* and not even potentially (access) conscious. In addition, the standard examples used as evidence for what I called the Integrative Nature of thought typically involve personal-level states of which we are introspectively aware and whose contents we express in natural language. So the fact that there are many sub-personal, unconscious cognitive states at least raises the question of whether every such cognitive

²⁴ For instance, as described by Marr 1982 chapter 1, and Cummins 2000.

state is integrated with the rest of the subject's knowledge in the way that thoughts are. Thus, there are cognitive states and processes that reach beyond the faculty of thought.

For our purposes, I propose we just say that cognition consists in all those mental (and thus intentional/representational/information-carrying) systems and mechanisms whose primary, evolutionarily developed functions subserve, enable, or are otherwise closely related to the exercise of thought—and leave it at that. This vague characterization leaves ample room for cognitive science to interpret it as they think best and to enumerate these systems and outline their boundaries for us.²⁵ And it allows dispositional states, and sub-personal representations and processes, to count as cognitive, alongside instances of genuine thought. (An analogous distinction can be drawn between perception in a narrower, everyday sense, and the various perceptual systems posited by science.) But the question about cognitive phenomenology can still be stated merely in terms of phenomenology that is proprietary to *thought*, in the everyday sense (note that both the Modal and Qualitative theses are formulated in terms of 'thought', not 'cognition'), since this is what parties to the debate were primarily interested in all along, as evidenced by the lists from Smithies, Chudnoff, and Bayne quoted above.

This way of distinguishing between thought and cognition has the distinct advantage of allowing us to move forward on the CP debate while setting to one side the extremely difficult and hotly contested question of how to draw a boundary between cognition and perception. Whatever the ultimate outcome of that investigation—even if the science at best only supports a very fuzzy line between the two—there still should be a

²⁵ This way of looking at the issues may seem to put me in tension with the likes of Masrour 2011, 27-30. But I am not absolving empirically informed philosophy of the responsibility of contributing to these and related debates. I am simply suggesting we can bracket the nitty-gritty of the cognition/perception issue for most or all of our treatment of the cognitive phenomenology question.

fairly robust divide between the commonsense conception of thought and thinking I have outlined above and our folk conception of perception. And this distinction is all we need to get our investigation of cognitive phenomenology up and running.

3. Thought and Concepts

There is a further conceptual-terminological issue that may have a direct bearing on the CP question, and which we cannot leave to cognitive science to sort out: the nature of *concepts*, or what we should take the term ‘concept’ to mean. For simplicity’s sake, I am going to stipulate what I take a concept to be. Concepts are all and only those mental representations that serve as the (semantically evaluable) constituents of thoughts. Thus, concepts, like thoughts, are concrete entities realized in the mind or brain, not abstract entities such as “modes of presentation” or the contents of these mental representations.²⁶

The idea that concepts are in some way closely tied to thought and cognition is hardly novel or controversial, but it is controversial to claim that concepts are *only* constituents of thoughts. However, the claim should be understood as referring to representation *types* rather than *tokens*: each concept type must be such that tokens of that type can serve as constituents of thoughts for the subject, and every constituent of an individual thought token counts as a concept. This requirement leaves wide open the question of whether or not a concept can be tokened without forming part of an individual thought.

Of course, putting things this way raises the tricky question of how we should type representations. On some ways of doing so, two representations might count as being of

²⁶ For discussion, see Fodor 2000, Chapter 3; Laurence and Margolis 1999, section 1; Peacocke 1992, 1996.

the same type, and thus as tokens of the same concept, even though one is part of a thought and another is not. On other ways of doing so, if two representations are tokens of the same concept, they will both be parts of thoughts.²⁷ But these different alternatives are not settled for us in advance simply by our understanding concepts as all and only those representations that serve as the constituents of thought.

Thus, I am suggesting we anchor our conception of concepts in the conception of thought we already possess. In this way we leave our initial characterization of what it is to be a concept relatively sparse, rather than weighing it down with additional details. Let me motivate this approach by comparing it to a possible alternative that adds to our definition of concepts some additional features they are supposed to possess. Jerry Fodor (2007) has argued that the phenomenon of “perceiving as”—e.g. seeing the T-Rex *as* a T-Rex or hearing the police siren *as* a police siren—always requires the exercise of conceptual capacities, which consists in the application of “discursive,” i.e. word-like, representations. As he puts it, “*seeing as ... and describing as...*, like other acts of conceptualization, operate by subsuming distal things under the concept that is expressed by the predicate of some

²⁷ If what individuates representation types is simply their content, then, absent an argument to the contrary, it seems possible that there could be tokens of the same representational type, and thus tokens of the same concept, some of which are constituents of thoughts and others of which are not, e.g. concepts employed in perceptual states that fail to count as thoughts by our criteria. But one might think that having the same content is a necessary, but not a sufficient, condition for two mental representations to count as the same type. So we might hold that representation types should be individuated by content plus inferential role, in which case the possibility is left open that two representations with the exact same content but different inferential roles (supposing that inferential role could differ between representations with the same content, e.g. between one representation used in perception and another used in thought) could count as belonging to distinct representational types. In this scenario, a subject’s perceptual faculties could have proprietary representations that are not concepts, but that represent some of the same things that our concepts represent. A third alternative is for representations to be typed by their content plus their intrinsic, “syntactic” properties, in which case there is at least logical space for a view according to which two representations with the same content and same syntactic properties, but with different inferential roles, nonetheless could count as tokens of the same concept, even though not every representation with that content would automatically count as that same concept. And again, in this case it could be that some tokens of that concept were constitutive of thoughts, while others were not.

mental representation. It is entirely in the spirit of RTM [i.e. The Representational Theory of Mind] that ‘conceptualizing’ and ‘predicating’ are two ways of talking about much the same thing” (110). He has recently taken Tyler Burge (2010) to task for not having a robust enough notion of perceiving-as, and consequently not fully appreciating the role that concepts must play in perception, even the perceptual capacities of infants and primitive animals (Fodor MS). Ned Block has responded by coming to Burge’s defense (2014).

On Burge’s view, perceptually representing, e.g., biological kinds such as *predator* or *conspecific*, only requires that the creature’s perceptual system possesses the appropriate “perceptual attributive”: a perceptually innate, non-conceptual type of representation that is predicate-like insofar as it is dedicated to attributing a kind, property, or relation to a perceived individual, but tokens of which, unlike predicates, necessarily are tied to a given perceptual context and to playing a role in fixing singular reference within that context (537-544).²⁸ Since perceptual representation involves “attributing” the kind or property in this way, perceiving *as* (falling under the kind or possessing the property) is accounted for in terms of perceptually representing by employing the relevant perceptual attributive.

Now, I am not interested in taking a side in this dispute, except on the meta-level question as to which use of the term ‘concept’ is most helpful to adopt in relation to the question of cognitive phenomenology. Considered in this light, Fodor’s understanding of conceptual representation, as equivalent to representation-as, is less than helpful. For,

²⁸ Burge also suggests that perceptual representation may be in a pictorial, or “map-like” format, rather than consisting of discursive representations, and it is crucial to his story of the context-dependence that distinguishes perceptual attributives from concepts that the former are *not* able to be integrated into a “propositional structure” (543). This point echoes what I earlier called the second “strand” of the Integrative Nature of thought: the general and compositional nature of thought constituents (i.e. concepts).

when combined with an understanding of concepts as thought constituents, it threatens to elide the distinct categories of thought (or cognition) and perception, and therefore any difference between cognitive experience and perceptual experience. If concepts, first and foremost, are what we think with, but they are also ubiquitous in the mind and in particular in perception, popping up wherever there is representation-as, then we might be led to conclude that thought-like representation invades perception to such a degree that there is no meaningful distinction to be made between perceptual experience and cognitive experience, or between phenomenology of perception and phenomenology of thought.

Of course, we can avoid this consequence and hold on to Fodor's view if we give up on defining concepts in terms of their being the constituents of thoughts. But then, to avoid confusion, we would need to stress that our use of the term 'concept' does not presuppose any special association between concepts and thought, which would amount to our adopting what I take to be a fairly unorthodox use of the term. Ultimately, everyone is free to use their theoretical terms as they see fit. But for our purposes, I prefer to reserve the category of conceptual representation for only those mental representations that can be part of thoughts, where, at the very least, the category of thought is meant to *contrast* with the category of perception.²⁹

Thus, I am siding with Block and Burge insofar as I am allowing that the perceptual representation of kind and other "higher-level" properties are not necessarily conceptual, though I do not think this point is necessarily at odds with Fodor's insight about

²⁹ It *could* turn out that perception is entirely non-conceptual even in Fodor's sense of 'concept'. But this result is implausible, since it would entail that all representation-as is post-perceptual, and in any event, we certainly cannot assume this view from the outset simply in order to separate perception from thought. If we wish to preserve this distinction for the purposes of asking about cognitive phenomenology, then our way of drawing the distinction better not be so vulnerable to collapsing.

perceiving-as requiring that a subject apply a discursive representation to an individual. For all I know, perceiving the T-Rex as a T-Rex does require applying the T-Rex representation—concept or perceptual attributive or whatever—to the perceived individual, and this representation is more accurately understood as being word-like rather than picture-like. I just am not prepared to call all such discursive representations ‘conceptual’ unless they can serve as the constituents of thought. In this way, we can continue to associate concepts with thoughts, without that association having the consequence that every experience of perceiving-as becomes potential evidence for a proprietary phenomenology of thought.

In other words, a virtue of this way of thinking about concepts is that it preserves the independence of the CP debate from other related philosophical debates in the literature. For instance, consider the contemporary debate about whether higher-level properties, such as affordances or natural kind properties, form part of the content of perceptual experience.³⁰ On my understanding of what concepts are and how they relate to cognitive phenomenology, granting that, e.g., kind properties are represented in visual experience does not necessarily (but may) lead to accepting CP. For, suppose we grant that representations of such kinds are an essential, constitutive part of a given experience, such that removing these representations would necessarily alter what it’s like to undergo the experience. This fact alone would not entail (a) that the experience is essentially conceptual—i.e. constitutively involves representational types, tokens of which serve as constituents of thoughts—nor (b) that if it is conceptual, *this very experience* essentially involves the operation of the faculty of thought. But for the phenomenology of the

³⁰ See Siegel 2006a, Pautz 2009, Bayne 2009, Masrour 2011, Nanay 2011, 2012.

experience to count as CP (on the Modal Thesis), it must token a particular thought, which necessitates the exercise of the faculty of thought.³¹

If a given perceptual experience with conceptual content were to necessitate the thinking of a thought, then the experience would be an instance of cognitive phenomenology (at least in the Modal sense). To take a famous example: the fact that seeing the duck/rabbit as a rabbit makes an essential phenomenal difference to conscious experience is evidence for CP just in case it is agreed that the relevant representation (assuming it is indeed a rabbit representation) does not belong to the visual system alone, but really is the concept RABBIT, the same one I use to think thoughts about rabbits, and further, that the use of that concept in this case constitutes an exercise of thought.³² Nothing in the way I propose we understand concepts, cognition, or perception settles these issues in advance.

Similarly, my way of construing these issues does not commit us one way or the other on the long-standing debate over whether, or in what ways, perception has conceptual or non-conceptual content. If a representation that serves as a constituent of thoughts makes a constitutive difference to a perceptual experience, then we should say that perceptual experience has conceptual content. This scenario also may be sufficient to

³¹ On the Qualitative Thesis, the issue is more flexible, since what we say here depends on how we interpret the requirement that the phenomenology in question be “specially associated” with thought.

³² Note that the fact alone that the representation in question is a concept is not sufficient to establish the existence of CP as part of the experience of the duck/rabbit. We would also need to show that the contribution made to the phenomenology of the experience by that concept could not be part of the experience in the absence of the deployment of the concept (cf. Carruthers and Viillet 2011). However, some philosophers, both Liberals and Conservatives, seem to think that representation of higher-level properties in perception that makes a constitutive difference (in the way just described) to experience suffices for cognitive phenomenology (for instance, see Carruthers and Viillet 2011, Strawson 2011, Montague forthcoming). But I think this is a mistake, for the sorts of reasons raised in the main text. It is noteworthy that Susanna Siegel, one of the foremost advocates of the higher-level content view concerning perceptual experience, explicitly distinguishes her view from acceptance of CP (2006a).

guarantee that there is cognitive phenomenology, but the two issues can come apart in several ways. First, there could be instances of CP that are not instances of perceptual experience with conceptual content, e.g. if there are experiences of “pure thought” consisting entirely of CP properties. Second, some philosophers who accept that perception is conceptual might deny that perceptual *experience* is. Third, some philosophers who accept that perceptual experience is conceptual might deny that this means the conceptual content must make an *essential* or *constitutive*, rather than merely causal, difference to the phenomenology of such an experience, in which case there could be instances of conceptualized perceptual experience that are not instances of CP. Fourth, we might accept that perceptual experience is conceptual but deny that this means it involves genuine thought, since it fails to exhibit one or more of our three main characteristics of thought set out earlier. Finally, many conceptualists about perceptual experience might have a view of concepts that is more in line with Fodor’s view, according to which to say that perceptual experience is conceptual is really only to claim that it is a species of representation-as and/or word-like representation. And this claim is clearly independent from an endorsement of cognitive phenomenology.

V. Road Map of Where We are Going

The thesis is divided into five chapters. **Chapter Two** performs some necessary ground clearing, by responding to a recent argument that thoughts cannot possibly be conscious experiences at all, since they lack the requisite metaphysical “profile”: they persist through time as states or events, rather than as processes. In response, I show how this argument does not threaten the possibility of CP.

In **Chapter Three** I evaluate the two major arguments for the existence of CP, the Phenomenal Contrast Argument and the Self-Knowledge Argument, and I respond to the positive case for the Modal Thesis. Part of my response involves comparing the Strong CP position with the view that perceptual experience represents higher-level perceptual properties. I formulate, and refute, a version of the Phenomenal Contrast Argument based on this “higher-level content” view. I conclude that, while the Liberal arguments offer limited evidence in favor of CP in the Qualitative sense, the evidence does not specifically support the existence of the type of CP required to determine the contents of our thoughts, which is to say it does not support the Modal Thesis. I further argue that the fact that Phenomenal Liberals often take their Phenomenal Contrast Arguments to support the Modal Thesis is best explained by their (often implicit) commitment to a certain contentious view of self-knowledge that is also crucial to the Self-Knowledge Argument, but which the Conservative will probably reject.

Chapter Four examines evidence for a more modest type of CP, inspired by William James’s notion of “fringe” consciousness. For example, we have all vividly experienced recognizing a face in a crowd as familiar without being able to place it, or sensing a word on the tip of our tongue. I argue there are compelling reasons to think that the phenomenology of these sorts of experiences is not reducible to the sensory, and is rightly considered cognitive, in that it serves to communicate information to the subject *about* her (largely unconscious) cognitive life. I then apply this account of Moderate CP to explain various contrast cases introduced in Chapter Three.

Finally, in **Chapter Five**, I lay out my main positive argument for the conclusion that phenomenology cannot determine thought. The crucial insight is that the Modal Thesis

depends for its intuitive plausibility on the notion that what is *thought* can be wholly captured in what is *experienced*—that of which we are aware in the special, phenomenal sense. For as I argued at the beginning of this chapter, the phenomenal features of an experience are presented to the subject, and thus are features of what the subject is directly aware of. But, as I will substantiate in detail, this view of conscious thought is not very promising. Indeed, it runs afoul of any realistic theory of how the conceptual contents of our thoughts are structured and individuated. Thus, our only satisfactory option is to reject the Modal Thesis, and with it Strong CP.

Chapter 2

Are Thoughts Ever Experiences?

I. Introduction

Near the end of the last chapter, we examined the relationship between thought and the categories of cognition and perception, and between thought and concepts, with the aim of clarifying how the ways we construe these relationships might impact on the debate concerning cognitive phenomenology. In this chapter, we will take another look at the nature of thought, by considering whether thought's *temporal profile*, how a thought

persists through time, can tell us anything decisive about the relationship between thoughts and *experiences*. Can thoughts literally *be* experiences, or at best are they only ever *accompanied* by certain experiences with phenomenal character?

This issue is of pressing concern for our study of CP, because Phenomenal Conservatives have recently defended a particular view of the temporal profile of thought and used it to formulate an ambitious argument against CP. The argument purports to show that the existence of CP in all forms is impossible, thereby rendering any further exploration of these issues futile. If the argument is sound, then the CP debate is settled, and this ought to be a very short thesis! Thus, the question of whether the argument is sound demands serious attention from all parties to the CP debate.

The argument in question belongs to Michael Tye and Briggs Wright (2011). Not content to explain away the apparent evidence for CP, Tye and Wright attempt to render the Liberal position untenable by establishing that thoughts are not the sorts of things that *could* have the requisite phenomenology. However, a careful assessment of their argument will reveal that it leaves cognitive phenomenology unscathed. In this chapter I will present their argument and related issues, and then I will offer two responses.

II. The Process Argument

Tye and Wright (henceforth T&W) present what I will call the Process Argument (PA) against cognitive phenomenology. Their exposition of the argument is extremely brief and relies on the work of Matthew Soteriou (2007), who in turn draws on observations made by P.T. Geach half a century ago. The core idea is that thought cannot be part of the stream of consciousness, because all items in a subject's stream of consciousness are

processive, i.e. they persist by “unfolding” as a series of events over time, whereas thoughts are *non-processive*, static states.³³

Even if we accept that a thought’s content has a complex structure, we should not make the mistake of supposing that elements of that structure “occur successively, as the words do in a sentence” (Geach 1957, 104). Nor should we suppose that the thinking of the thought lasts as long as it takes to say the sentence in inner speech (Anscombe and Geach 1961). On the contrary, “once one begins to think that claret is delightful, one has already *achieved* the thinking of it...the whole thought arrives at once” (Tye & Wright 2011, 342). By contrast, mental items in the conscious stream “unfold over time in the way that an event like a cricket match unfolds”: by having distinct temporal parts at each successive moment of their existence (342). One might capture this difference by saying that cognitive states, unlike cricket matches and conscious experiences, are “wholly present” at each moment they persist.

The relevant distinction is *not* that processes have temporal parts while states do not: we can conceive of “slicing up” enduring states into shorter temporal segments. Rather, the difference amounts to the fact that each of these temporal parts, considered on their own, already counts as that very state (type), while the proper temporal parts of a process cannot be identified with the process (type) itself. For instance, I am in a state of anxiety all day because I am in this state at each moment between sunup and sundown, whereas I am in the process of writing this chapter all day by enacting successive parts of

³³ The relevant sense of the term ‘thoughts’, here, seems to be (a) judgments and related propositional attitudes, and/or (b) the mental representations that represent the propositions towards which we bear the various attitudes. This is a narrower use of the term than the one I introduced in the previous chapter, since there I allowed that the content of individual thoughts could be sub-propositional. But the success or otherwise of the argument does not seem to hinge on this difference, since the strategy I rely on to refute the argument requires that I focus on propositional thoughts as well.

this process throughout the day, *parts that do not independently count as the whole process*. So a state's identity is already determined at the level of each individual temporal segment, but not so with a process, whose identity depends constitutively on all its temporal parts.³⁴

T&W acknowledge that inner speech and imagery might constitute one's "experience of thinking," but they claim such experience must be distinguished from the thought itself. Since thoughts are non-processive states, it is a mistake to think they could ever show up in one's stream of consciousness. But the domain of phenomenal consciousness just is that which makes up the stream of consciousness. Therefore, thoughts "simply aren't suited to be the bearers" of phenomenology, and there is no such thing as a phenomenology of thought (343).

Here is my best attempt at capturing T&W's argument:

Process Argument (PA):

- P1** Thoughts do not unfold as processes
- P2** Everything that is part of the stream of consciousness unfolds as a process
- P3** All experiences are part of the stream of consciousness
- C** Thoughts are not experiences

I say a thing "unfolds as a process" as shorthand for a thing persisting through time by having a processive structure, or else being the type of thing that would have a processive structure if it were to persist.³⁵ The above formulation is valid. And it captures T&W's

³⁴ See Chudnoff 2015, chapter 4; Velleman & Hofweber 2011, 16; and Crane 2013, 14-16. Many philosophers in conversation seem skeptical of the distinction being drawn between states and processes. I am granting that there is something to the notion for the sake of exploring my opponent's argument.

³⁵ This qualification is needed in order to rule out the idea that an instantaneous thought could also be an instantaneous experience. Unqualified, the argument would be invalid: from premises about *persisting* thoughts and experiences, we could not reach the conclusion that *all* thoughts are not experiences. Another way to deal with this issue is to deny that there are any instantaneous experiences. On this alternative, even if there are instantaneous thoughts, which (perhaps) can be part of the conscious stream, they nevertheless

claim that thoughts cannot be “the bearers of the relevant phenomenology.” For, as we are using the term in this thesis, all and only experiences have phenomenal character, i.e. instantiate phenomenal properties. In other words, T&W’s argument is an attack on the idea that thoughts can be the *vehicles* of phenomenal character.

In trying to respond to this argument in defense of CP, we have two options. We can challenge one or more of the premises, or we can accept the premises but deny that the conclusion undermines the Liberal position. In what follows, I explore these options in turn.

III. The Simple Response

1. Suppose we grant all three premises of the Process Argument. Why should we think its conclusion amounts to, or entails, the claim that CP does not exist? The answer depends on how the claim that thoughts are not experiences relates to whether there is cognitive phenomenology in either the Modal or Qualitative sense. The Qualitative Thesis merely posits a type of phenomenology that is distinct from sensory phenomenal properties and that is somehow specially *associated* with cognition. While according to the Modal Thesis, there exists phenomenology of the following type: there is no metaphysically possible scenario in which a subject is enjoying an experience with this phenomenology without also entertaining some thought or other. But I see no reason why these cognitive phenomenal properties must be properties *of* the thought they necessitate, such that the thought itself counts as an experience in the stream of consciousness.

cannot be experiences, since all experiences are persisting parts of the stream. There is some textual evidence for thinking that Soteriou, at least, believes that all experiences have duration (2007, 562). O’Shaughnessy’s “Absolute 0” argument (2003, 42ff) might be taken to imply this claim, too.

One might object that we have mischaracterized the relevant CP thesis: the Process Argument *is* incompatible with any claim that there is a type of phenomenology that is *instantiated* by a thought or cognitive state. But T&W have not provided a reason for thinking the Phenomenal Liberal must adopt such a thesis. Thoughts and experiences might be distinct mental entities; nonetheless, it might be that the phenomenology of certain experiences systematically relates to thought in important and interesting ways such that we are justified in labeling it “cognitive phenomenology.”

So, here is my first response to T&W’s argument: for all they have said, we are free to reject the assumption that if thoughts are not, strictly speaking, experiences then cognitive phenomenology cannot exist. The Liberal does not have to reject any of the premises of the argument, since the conclusion of the argument does not lead inevitably to a denial of CP. Call this the Simple Response to the Process Argument.

2. I will consider two objections to the Simple Response. First, if thoughts lie outside the stream of conscious experience, then why think of them as phenomenally conscious at all? And if thoughts are never phenomenally conscious, one might wonder whether it makes much sense to think of certain phenomenal properties as *proprietary* to thoughts. Consider the Modal Thesis: even if CP experiences somehow necessitate thoughts, should CP really count as phenomenology *of* thought if these thoughts are not rendered conscious by CP? The worry is just as pressing for the Qualitative Thesis: can we give a satisfactory account of what it takes for a CP property to be “specially” associated with thought if we deny that a CP property can make a thought conscious?

In response, we might start by distinguishing between experiences and phenomenally conscious states.³⁶ A phenomenally conscious state is any mental state that is associated with an experience such that there is something it is like for the subject to be in that state. There is often something it is like to be in a certain state (token), e.g. a particular attitude state, even though the state is not individuated by what it is like to be in it: one could be in a state of that same type with different phenomenal character or with none at all. It would seem to be an open question whether such a phenomenally conscious state must instantiate the phenomenal properties in virtue of which there is something it is like to be in it. Indeed, even T&W recognize a distinction between thoughts and the experiences of thinking those thoughts.³⁷ They must recognize, then, a legitimate sense in which the experience is an experience *of* thinking, even if, as they claim, the thought state itself does not enter the stream of consciousness. And presumably thoughts that are thus experienced thereby count as phenomenally conscious.

However, it seems we have merely shifted the problem, from specifying what it takes for a thought to be conscious, to specifying what it takes for there to be “something it’s like” to have a thought. When does an experience constitute what it’s like to be in a particular phenomenally conscious state? Clearly, the relation that holds between the experience and the thought must involve more than mere co-presence in the mind for some duration of time. Nor should we say that in order for an experience to count as being what it’s like to think a given thought, a subject’s having the thought (type) must be either necessary or sufficient for the subject to have the experience (type). However, it seems

³⁶ Cf. Bayne and Chalmers (2003).

³⁷ “What is it that unfolds over time if not the thought? One obvious suggestion is that it is that which accompanies the thought – those various phenomenal goings-on that one’s experience of having the thought comprises” (Tye & Wright, 342).

there does need to be a strong correlation in a subject's mental life between the experience type and the thought type. I don't pretend to have anything especially clever or original to offer on this point. One option is just to say that the experience of a thought is that part of the subject's conscious experience at a time that obtains *somehow* in virtue of the thought (token), or else that the thought obtains in virtue of the experience (token).

Ironically, Matthew Soteriou endorses such a picture of the relationship between phenomenal experiences and cognitive states, in the very paper to which Tye and Wright appeal in order to support their argument. Crucial to Soteriou's account is his observation that when we want to pick out mental states that exist in virtue of some phenomenally conscious event or process, we might do so "at different levels of abstraction" (561). For instance, we can pick out a cognitive state as the state that obtains in virtue of an experience with such-and-such phenomenal character, or alternatively, we can pick it out without making any reference to phenomenology, as the state with such-and-such intentional content. Similarly, we might pick out the phenomenal event or process itself at the level of phenomenal character or intentional content.

Since in general we wish to leave open the possibility that both cognitive states and mental events could be tokened with different phenomenology than they actually have—or could lack it altogether—we tend to refer to these states/events without mentioning their phenomenal natures, e.g. as the state of believing that *p*, or as the act of judging that *p*. Nevertheless, Soteriou contends, a mental event or process such as a judgment might possess phenomenal character and so be part of the stream of consciousness. And the enduring state that obtains in virtue of such an event can be said to be phenomenally

conscious in a secondary sense, by being the direct result of this conscious experience (560-562).³⁸

If we accept at least the spirit of Soteriou's account, we can respond to the Process Argument by accepting all the premises but denying that the conclusion entails the denial of cognitive phenomenology. For on this picture to deny that thoughts—conceived as occurrent cognitive *states*—are experiences does not mean that there cannot be phenomenal properties that characterize an experience of thinking, or that we cannot speak of phenomenally conscious thoughts. Furthermore, if some or all of these phenomenal properties are qualitatively distinct from sensory phenomenal properties, then the Qualitative Thesis is vindicated. And if some or all of these phenomenal properties *necessitate* the obtaining of thoughts, then the Modal Thesis is vindicated.

3. However, there is a second objection to the Simple Response, namely that it will only work for certain Liberal positions. For, as we will consider in detail in Chapter 5, some defenders of cognitive phenomenology go beyond the assertion of a bare necessary connection between CP properties and thoughts, to endorse an explanatory or metaphysical *dependence* of content on phenomenal character, in order to account for why

³⁸ Soteriou understands the relation between the experience of judging that *p* and the belief state that *p* as being constitutive (551-552), whereas I want to leave open the question of the specific nature of the “in virtue of” relation that holds between experience and thought (or vice versa). A word on terminology: unlike Soteriou, I prefer not to describe the occurrent cognitive state as a “belief,” as I explained in chapter 1, since I understand all beliefs to be dispositional. However, according to the sort of response to the Process Argument I am currently exploring, what I call a “judgment” or a “thought” corresponds *not* to the event/process in the conscious stream (what Soteriou calls an act of “judging”), but to the cognitive state that exists in virtue of this event/process. In what follows I explore the possibility of identifying thoughts/judgments with the entities in the phenomenal stream themselves.

the Modal Thesis is true in the first place.³⁹ But affirming one of these dependence relations threatens my response to PA.

Consider the claim that the Modal Thesis is true because CP properties are *identical* to the intentional properties of cognitive states: the properties that make up what it is like to entertain a given thought content are the very same properties that constitute the content entertained.⁴⁰ Proponents of CP sometimes say things that suggest this identity claim, which offers an elegantly simple explanation of why the Modal Thesis is true: if a thought's phenomenal properties are identical to its intentional properties, then for there to be a CP experience there would also have to be a thought with intentional content. However, if the identity claim is true, then cognitive phenomenal properties must be properties *of* the thought they make conscious, since surely the intentional properties of a thought are properties of it (after all, a thought is at least partly individuated by its content). But then, the identity claim is incompatible with the conclusion of PA, that thoughts are not experiences. So my Simple Response to the Process Argument would fail for any account of CP that included this identity claim.

A different way for thought content to be dependent on cognitive phenomenology is the somewhat vague suggestion that a specific content is "bound up with" and just "falls out of" the specific phenomenal character of each CP experience. On this view, the relation between content properties and phenomenal properties may not be identity, but it is something very close. An example of this sort of claim comes from David Chalmers (describing, with approval, related views of Colin McGinn and Charles Siewert):

³⁹ For instance, see Kriegel 2013; Horgan & Tienson 2002.

⁴⁰ Of course, it is not quite right to say that identity could be a relation of asymmetrical dependence between intentionality and phenomenology, but the identity claim is similar to the dependence claim in that it can be offered as an explanation for why the Modal Thesis is true.

...there is a strong intuition that something about representational content is *internal* to phenomenology. That is, given a specific phenomenology, it seems that if a mental state has this phenomenology, it must also have a certain specific representational content...the basic idea is that visual experiences are *assessible for accuracy*, in virtue of their phenomenal character. For example, when I have a visual experience as of something X-shaped in front of me, this experience may be either accurate or inaccurate, depending on what is really in front of me. Further, it seems that any visual experience with the same phenomenal character would be assessible for accuracy in the same sort of way. (2004, 6)

This passage describes perceptual experiences, but many proponents of cognitive phenomenology would apply this view about phenomenally grounded intentionality, or “phenomenal intentionality,” to the relationship between the phenomenology and content of our thoughts.

Call this sort of view a Phenomenal Intentionality (PI) claim. If Liberals accept the Modal Thesis because they accept a PI claim about the relationship between CP and thought contents, then their version of CP is vulnerable to the Process Argument. The reason for this vulnerability is not that the PI claim entails that the phenomenal properties in virtue of which a mental state has its content must be instantiated by that very state (though, admittedly, it would be strange if this were not the case). Rather, the problem arises because the PI claim implies that the subject’s entertaining of a thought content, in being fixed by the subject’s phenomenal experience as it is undergone, should *unfold* as the experience unfolds. The experience, *and therefore the thought*, is a process. In other words, the PI claim may not directly conflict with the conclusion of PA, but it does seem to conflict with PA’s premises.

Contrary to the Simple Response, a defender of this version of the Modal Thesis cannot accept the assumptions and reasoning of the Process Argument yet deny that the argument threatens the existence of CP. Thus, we have seen that this initial reply to the

Process Argument only works for defenders of the Qualitative Thesis and for some defenders of the Modal Thesis. I do not think the problems posed for other accounts of Strong CP are insurmountable. But in order to defuse them, we need to adopt a different response to Tye and Wright's argument, one that does not seek to make CP consistent with the conclusion of the argument, but instead rejects one or more of its premises.

IV. Are Thoughts Processes?

1. Let me remind us of the details of the Process Argument (PA), as I formulated it:

P1 Thoughts do not unfold as processes

P2 Everything that is part of the stream of consciousness unfolds as a process

P3 All experiences are part of the stream of consciousness

C Thoughts are not experiences

In what follows, I will direct my attack against Premise 1.⁴¹ Before proceeding, I need to make one qualification, and draw an important distinction.

First: the distinction. T&W, Soteriou, and Geach each motivate P1 with the observation that thoughts occur in the mind "all at once," which might be understood as an appeal to introspection. For instance, T&W assert, "Thinking the thought does not unfold in the way that a string of sounds from a piano unfolds in an etude. The whole thought arrives at once" (342). If they are attempting to describe how it seems to us from the first

⁴¹ I also believe Premise 2 is vulnerable to objections. One major problem for P2 is that it seems to require denying that in cases where a part of our total perceptual experience stays the same for some time (e.g. a persistent visual experience of a blue wall) the conscious stream is partially constituted by enduring phenomenal states (e.g. a phenomenal blue state). For discussion, see Chudnoff 2015, chapter 4; O'Shaugnessy 2000, 42 and 2003, 42ff; Oliver Rashbrook MS, 16-18; and Soteriou 2007, 547-550. But I am more interested here in exploring what we should make of the argument if we grant P2.

person perspective, then they are not ruling out *a priori* the apparent possibility that we think in unfolding stages.

However, Tye and Wright introduce their argument by saying, “we shall conclude by arguing for something stronger than the simple claim that in actual fact there *is* no such thing as a phenomenology of thought...We shall offer considerations regarding the nature of thought itself that suggest that thoughts *could* not be bearers of the relevant phenomenology” (341). For T&W thoughts “do not have the right structure” and so “simply aren’t suited” to be part of the stream of consciousness (341, 343). Now, it is open to them to contend that the conclusion of their argument, while metaphysically necessary, is arrived at by *a posteriori* investigation. But it is difficult to see how the case they present, if construed as built on a posteriori grounds, could be sufficient to establish that thoughts, by their very nature, *must* occur “all at once”.⁴² Nevertheless, someone could claim that, merely as a contingent fact about human cognitive architecture, thoughts are non-processive and thus non-experiential, and presumably this view would not be something we could establish from the armchair. Thus, the key distinction we need to make is between versions of the Process Argument that are based on **A Priori** considerations and those that are based on **A posteriori** considerations. I will consider—and reject—P1 in both versions of the argument.

Second: the qualification. Thoughts, understood on the model of judgments, constitutively involve the bearing of an attitude to a content. In challenging P1, I want to leave open the question of whether the bearing of any attitude beyond merely *entertaining*

⁴² Soteriou (2007, 544-547) and Geach (1969, 64) make much of *linguistic* evidence for their position. Although evidence of how we talk about thought is a posteriori, one can take such observations as corroborating evidence for a priori analyses of our concepts of thinking and thought. This second approach arguably characterizes Geach’s line of reasoning.

a content could be a process or not.⁴³ For instance, one might argue that for a subject to take up certain attitudes towards a content essentially involves accepting certain normative commitments, and one might think that standing in these normative relations is an all-or-nothing matter. So coming to bear such an attitude cannot be a gradual process, but must likewise be all-or-nothing. Regardless of what we take the merits of this argument to be, I am not going to address it in what follows. For it seems to me that it is possible in principle for the cognitive mechanisms that underlie the occurrent representation of a thought content to be independent from those that underlie the taking up of any number of more specific, “normative involving”, attitudes towards that content—e.g. endorsing it, doubting it, desiring it to be the case—such that the former could occur as a process even if the latter could not. My claim in what follows is only that the *representing of the content* may be a process and thus need not occur “all at once”. No doubt this is a substantive move, and it raises a number of deep issues, but fully addressing these issues lies beyond what I can accomplish in this paper.

2. Let us start by looking at the **A Priori** argument. The idea motivating P1 can be traced back from T&W to Soteriou to Geach, who writes, “unless the whole complex [of a thought] is grasped all together...the thought or judgment just does not exist at all” (1957, 104); and elsewhere: “unless the whole content of the thought that all tigers are dangerous is simultaneously present to the thinker, no such thought occurs at all.” The reason he gives for this is that “there is no such feat as thinking of all tigers except in the context of a

⁴³ I remain neutral about whether entertaining that *p* itself should count as a type of propositional attitude, or rather as something that forms the basis of every attitude state. If I adopt the former view, then I need to argue that this attitude does not bring with it the sorts of potentially problematic normative demands that other attitudes, such as belief, do.

thought that all tigers are so-and-so,” and “though I may no doubt have an indefinite thought of danger, a thought of tigers followed by such a thought would not be a thought that all tigers are dangerous” (1969, 34-35). We can borrow a thesis from Chudnoff (2015, 90), in order to capture Geach’s constraint on thinking:

Simultaneity: Suppose you think that *p* at *t*. Then all the parts of your thought that *p* occur at *t*. (43)

Geach’s reasoning is that (i) the logical-semantic relations holding between the various elements of the propositional content of a thought ensure a degree of unity in this content, and (ii) this unity requires the thought to be thought all at once, i.e. each part of the complex content must be “grasped” simultaneously. Since this unity appears to be a feature of all thoughts simply in virtue of their having propositional content, it seems Simultaneity can be justified by appeal to a priori reflection on the nature of thought alone. In a slogan: of a priori necessity, *one cannot think half a thought*. And this is a problem for the notion of thinking as a process, since, if Geach is right, a thought cannot be built up over time out of distinct (semantically evaluable) parts. So Geach’s Simultaneity principle provides a reason for accepting the A priori version of P1.

3. But the Geachian support for P1 does not stand up to closer scrutiny. Standard contemporary views of human cognition, the Representational Theory of Mind (RTM) and Language of Thought Hypothesis (LOT), raise problems for Geach’s ideas.⁴⁴ According to RTM, for a subject *S* to think a thought that *p* is to bear a (functionally specifiable, physically realized) relation to a representation in *S*’s mind that means that *p*. A thought is just a representation or symbol in the head, a “vehicle” with representational content. The

⁴⁴ For RTM and LOT, see Aydede 2010; Fodor 1975, 1987; Pitt 2013.

LOT thesis adds to this picture the idea that thoughts are *complex* representations, built up out of primitive representations according to a combinatorial syntax and semantics, such that “the semantic content of a molecular representation is a function of the semantic content of its atomic constituents together with its syntactic/formal structure”(Aydede 2010).

These popular theories are in serious tension with Geach’s views about thought. For starters, if to have a thought is just to token a certain kind of mental representation, and a mental representation is identical to, or realized by, a neural event, then having a thought boils down to a certain kind of neural event. But neural events arguably persist processively, by having distinct parts at each moment they exist. So it seems that thinking a thought also must unfold over time.

The defender of Geach can insist that a thought does not necessarily have the same temporal structure as any neural event that happens to bring it about. For one thing, familiar considerations about multiple realizability mean that the representation is at best realized by, but not identical to, a neural process in the brain. And for another thing, we are free to conceive of a thought representation as realized by the neural event/state that exists only at the very end of an extended neural process, such that the entire process is not itself necessary for the tokening to take place. In other words, while various cognitive activities might rightly be considered processes, it would be a mistake to consider the thought token that *results* from these activities to be a process. Instead the thought, on this view, is identical to the instantaneous event of the cognitive process being *completed*.

However, if this is the best Geachian reply to the above worry, then it runs into trouble when we add an LOT framework to the RTM. According to LOT, thinking a thought

literally involves the tokening of its constituents, and the propositional content of the thought is built up from the contents of these simpler representations. But then, RTM and LOT seem to leave open the possibility that the overall representation of a thought content could occur *non-simultaneously*, since the concepts that jointly constitute the thought representation could be realized in the brain consecutively.

It is no longer a viable option to identify the thought with an instantaneous event. For if thinking the thought *I dream of Jeannie* involves tokening the complex representation A-B-C, and this representation has as parts the concepts A, B, and C—meaning *I*, *to dream of*, and *Jeannie*, respectively—and further, if tokening A-B-C in fact involves tokening first A at t_1 , then B at t_2 , and finally C at t_3 , then we cannot identify the thought with the cognitive state only realized at t_3 (or with some further resulting state that only obtains after t_3). For LOT's commitment to the idea that thoughts have a constituent structure entails that the tokening of each constituent is literally part of the tokening of the whole. In this case, it *would* be possible to “think half a thought”—e.g. by tokening A and B but not yet C—before getting interrupted.⁴⁵

To be clear: I am not claiming that the LOT must work in our minds in the way that I just described, with thought representations instantiated in successive phases. Rather, my claim is that LOT must allow for the *possibility* that thoughts are tokened in our minds in this way, given LOT's general account of what a thought is. And since Geach's Simultaneity principle is intended as a universal constraint on all possible thought, we cannot accept

⁴⁵ One might object that if one were to get interrupted, the tokens of A and B would not necessarily count as half a thought, but simply as unconnected tokens of the concepts *I* and *to dream of*. We can grant this point for the sake of argument (whether it is right or not depends on when and how the syntactic properties that unify the thought are instantiated—see below). But it doesn't change the fact that, if thinking a thought essentially involves tokening a string of non-simultaneous simple concepts, then the thought cannot be an instantaneous event. Nor can the thought be an enduring state, since the temporal parts of the thought are not themselves that same (or indeed any) thought.

both LOT and the Geachian view.

4. An obvious reply to my argument in the last section is “well, so much the worse for these empirical theories.” If PA, conceived as an a priori argument, is incompatible with RTM and LOT, why should that be a reason to reject it? If these empirical theories really are incompatible with Geach’s insights about the nature of thought, why not just think Geach has hit upon an a priori reason that they cannot be true?

I don’t think this is a promising response. To see why, consider how LOT reveals what is deeply wrong with Geach’s view in the first place. From Geach’s argument we formulated the following principle:

Simultaneity: Suppose you think that *p* at *t*. Then all the parts of your thought that *p* occur at *t*. (Chudnoff 2015, 90)

What is striking is that Geach’s motivation for Simultaneity is very similar to things that Fodor says *in defense of* a “classical,” i.e. LOT, cognitive architecture in his dispute with connectionism: we need thoughts to be representations that are compositional, built up out of simpler representations according to syntactic rules, precisely because representing a proposition is *more than* just representing its semantic parts.⁴⁶ To think the thought *I love Lucy* requires something over and above the process of first thinking *I*, then thinking *love*, and finally thinking *Lucy*. As Geach puts it, “unless the whole complex is grasped all together” there just isn’t a single thought there at all. This is essentially principle (i) in my summary of Geach’s reasoning above.

But in Fodor’s characterization of the language of thought, LOT’s method of unifying individual concepts into a single thought with propositional content does not appeal to

⁴⁶ See Fodor & Pylyshyn 1988 for a classic elaboration of this idea.

temporal unity at all—rather the unity is *syntactic*. According to Fodor, the difference between a subject tokening a string of merely associated concepts and thinking a thought lies in the structural-functional relations between the different conceptual representations constituting a genuine thought, relations that encode the combinatorial structure of a representation in the language of thought.⁴⁷ In the case of a subject merely tokening a group of concepts, by contrast, this syntactic structure is missing. But there is nothing about LOT, as far as I can see, that requires the entire syntactic structure of a complex mental representation, along with the atomic representations that it combines, to be realized in the brain at a single moment. My suggestion is that the binding together of the various constituent representations into a single thought *could itself be a process*, unfolding alongside the representations. For instance, in the case of the thought *Ricky loves Lucy*, already at t_1 the mind is processing the concept of Ricky *as a part of an as-yet-unfinished propositional representation*. So it is appropriate to describe what goes on at t_1 as the beginning of a thought, the beginning of the entertaining of a proposition about Ricky, rather than just the deployment of the Ricky concept.

One might wish to quibble with details of Fodor's account, but the point is that Geach's central insight—concerning the difference between mentally representing a proposition and merely representing its semantic constituents—can be acknowledged, and accounted for, by RTM in terms of LOT. Indeed, making sense of this difference is a major motivation for positing a language of thought. But then Geach's insight falls short of providing justification for the Simultaneity principle, because it only gets us as far as justification for accepting (i), but does nothing to justify the further assumption, (ii). For

⁴⁷ See fn. 9 in Fodor and Pylyshyn 1988, 13-14.

being “grasped all together” does not require being grasped *in a single instant*, but only being grasped *as a unity*. And grasping a complex representation as a unity might involve grasping its parts in succession, but grasping them *as related* to one another. Thus, Geach and Fodor should be able to agree on the following alternative to Simultaneity:

Unity: Suppose you think that *p* at *t*. Then all the (semantically evaluable) parts of your thought must be structurally unified, where this means, at a minimum, (a) they must stand in certain relations, *r*, to each other, such that the tokening of these *r*-related parts is jointly sufficient for the subject to entertain the proposition, *p* (whether or not each part is tokened at *t*); and (b) merely tokening these parts of the thought successively or simultaneously⁴⁸ is *not* sufficient for the subject to entertain *p*.

According to LOT, the set of *r* relations consists of functional properties physically realized in the brain, which constitute the combinatorial structure of a thought. But one can accept Unity without accepting LOT, by giving an alternative account of the *r* relations that bind the parts of a thought together into a proposition-expressing whole. In any event, the Simultaneity principle—and the A priori version of PA that it undergirds—has little left to recommend it.

5. At this juncture, my opponent can try retreating to a more concessive position while still holding on to Premise 1. Perhaps, although a thought token could be constructed in successive stages as suggested above, there is nevertheless some *additional* feature of the nature of our thoughts that requires them to have a state- or event-like, rather than processive, structure. Perhaps it is the subject’s conscious *awareness* of the thought that cannot occur in stages.

We can put this idea by distinguishing between merely representing a thought

⁴⁸ Cf. Fodor and Pylyshyn 1988, 24-27.

content, on the one hand, and consciously entertaining or grasping a thought, on the other. The proposal is that consciously “grasping” a thought in this sense—i.e. becoming non-inferentially, first-personally aware of its propositional content—happens instantaneously if it happens at all. Indeed, my opponent might insist that the debate was over this sort of personal-level awareness all along: thinking, in the relevant sense, just is becoming aware of (a series of) thought contents. From this perspective, the key claim that leads to endorsement of P1 is the following:

Simultaneity*: Suppose you are consciously entertaining p (i.e. enjoying direct, first-personal awareness of p) at t. Then you are consciously entertaining the entirety of p (i.e. are aware of all the constituents of p) at t.

I think the prospects of finding adequate a priori support for Simultaneity* are dim. Remember, the sort of conscious awareness of thought content that figures in Simultaneity* cannot be awareness brought about by *phenomenal* consciousness. But if not, what sort of account of the awareness of our thoughts is left to us except one that ultimately invites a functionalist characterization? Certainly access consciousness (Block 1995) and related notions are functionalist notions. And if we understand this awareness in functionalist, representationalist terms, then we are back where we ended up in our discussion above: *a subject’s awareness or conscious grasping of a thought might be a neurally-realized functional process just as much as any mere representation of the thought content might be.*⁴⁹ If so, the remarks I made earlier about the possibility of the unity of a thought being achieved syntactically and non-simultaneously will apply *mutatis mutandis*

⁴⁹ This point holds whether we adopt a first-order or higher-order functionalist account of awareness of our thoughts. Cf. Joe Levine’s (2011) distinction between “implicit self-knowledge” and “explicit self-knowledge.”

to the unity of the subjective awareness of a thought.⁵⁰

6. We can turn, finally, to the **A Posteriori** version of the argument. What sorts of a posteriori considerations can we bring to bear in support of Premise 1? The options would appear to be (1) introspection on our own thoughts, and/or (2) what science can tell us about cognition. As for (2), if there is compelling scientific evidence that the LOT hypothesis is false, I have not heard it. Indeed, the evidence would also have to rule out any *prima facie* plausible account of thought that is consistent with the possibility of a thought unfolding as a process, and that meets the constraint placed on thought by Unity without relying on Simultaneity to do so. If one knows of such evidence, answers on a postcard, please. In the meantime, we are left with (1), our introspective reports.

Suppose the proponent of Simultaneity (or Simultaneity*) claims that we introspectively seem to entertain a propositional thought at a single moment. How compelling is her claim? Is the claim that it *always* seems this way, or only on the handful of occasions she happened to attend closely to her thoughts? It is not at all obvious to me that I can “look inward” and pinpoint the exact moment a thought begins and ends. Indeed, there are times when it seems fairly intuitive to describe a single thought as having developed during the thinking of it. An obvious example is thoughts caused by reading written language: I read ‘Ricky is’ at the bottom of one page, take a moment to turn the

⁵⁰ On this possibility, the awareness relation between a subject, S, and, e.g., the proposition that Ricky loves Lucy would be instantiated for the period of time it takes to become aware of the individual conceptual elements of the proposition (and how they relate to each other). At time t_1 , the instant that S entertains *Ricky*, S would count as entertaining a proposition (in the sense of Simultaneity*) derivatively, i.e. in virtue of undergoing an event that is partly constitutive of entertaining a proposition, just as long as the other events that are constitutive of this process occur at t_2 , t_3 , and so on. By analogy: I count as reading a sentence at a given moment in virtue of reading a word at that moment—an event that is partly constitutive of reading the whole sentence—and not in virtue of reading the whole sentence at that moment.

page (the pages are stuck together), and then read ‘in love with Lucy,’ at the top of the next page. I understood the first part of the sentence, but had no idea how it was going to finish until I turned the page. In order to say that the thought caused by reading this sentence does not unfold over time, we must insist that the first part is still being entertained when I arrive at the end of the sentence, so that the whole thing can be grasped at once.⁵¹ But must it *always* be this way? And is this question really something we can settle just by introspecting? Of course, the cognitive effects of reading and understanding the first part will still be unfolding when I finish the sentence, but that is different from the stronger claim needed here, that I am still somehow *holding in mind* the first part of the sentence at the end. Any temptation to insist that we always entertain the whole of the meaning of a sentence at the exact moment we finish reading it—no matter how long the sentence—likely originates from the lingering influence of Geach’s a priori reasoning about Simultaneity, which we have already seen is flawed.^{52 53}

However, for the sake of argument, suppose we accept the claim that we seem to be aware of our thoughts all at once. Still, we can distinguish between seeming to be aware and actually being aware, and accept the one while suspending judgment on the other. My

⁵¹ Alternatively, one might argue that even at the beginning of the sentence, I have already grasped a whole thought if I have understood anything at all, since I do not understand ‘Ricky’ as meaning just *Ricky*, but always *Ricky is something-or-other*—a full proposition, which gets altered as the details of the rest of the sentence get filled in. So rather than supposing the thought is an instantaneous event occurring after I finish reading, on this view the thought is an enduring static state (or, there are multiple thoughts corresponding to a series of such states), which begins when I read the first word. But the idea that every time I read, e.g., a single name, I have already grasped an entire proposition seems implausible. As in the main text, I suspect that any plausibility this view enjoys comes from the Geachian conception of thoughts, rather than directly from introspection.

⁵² A long sentence may represent a number of conjoined simpler propositions, in which case the subject would only need to be entertaining the final simple proposition at the end of reading, not the entire meaning of the sentence. But the general point still stands.

⁵³ Even if introspection did consistently suggest to us that thoughts happen all at once, it is unclear how this kind of evidence could justify the ambitious claim that T&W make about their argument: “We shall offer considerations regarding the nature of thought itself that suggest that thoughts *could* not be bearers of the relevant phenomenology” (341).

opponent might object that to seem to be aware of a particular proposition implies that one is aware of it, at least in the episode of “seeming” itself. In response, we can grant that seeming awareness of *p* and awareness of *p* cannot come apart entirely. We only need to argue for their coming apart in a more modest way: that it is possible for an awareness of *p* to seem instantaneous when in fact it takes time, and for instances of awareness of the constituents of *p* to seem simultaneous when in fact they are not.

The distinction between the actual temporal properties of events and the temporal properties they subjectively seem to have is a familiar one to philosophers of mind. Why might a subject be unaware of the true temporal relations holding between the parts of her occurrent thought?⁵⁴ We can ask this question with respect to cognitive awareness of a basic sort—the awareness we have of our thoughts simply by being awake, attentive, and in a lucid frame of mind—or with respect to introspective awareness that comes via deliberately attending to, or reflecting upon, our thoughts. Roughly, the former is awareness of *what* we are thinking—just the content—while the latter is awareness *that*, and perhaps *how*, we are thinking it. It is natural to suppose that explicit sensitivity to features of our thought tokens and acts of thinking themselves is a sophisticated capacity that we only tap into, if at all, on those occasions when we deliberately “turn inward.” So the basic awareness of a thought content might occur in stages, and yet we should expect a subject not to notice this fact unless she introspects.

On the other hand, it would not be surprising if we fail to notice the way our awareness of a thought content is built up out of stages even when we do introspect. For plausibly an intentional state makes us aware of just its content, and arguably the content

⁵⁴ Alternatively: between the parts of the (first-order or higher-order) state/process of awareness of a thought.

of an introspective state about a thought is simply that we are thinking a thought with a given content. So introspection might give us no information about *how* we are thinking the thought. Furthermore, even supposing we can learn about the temporal properties of our thoughts from introspection, the time differences between when distinct parts of the content are grasped might lie below the threshold of what is discriminable for a normal subject.

Finally, even allowing that differences in time between the stages of a thought are in principle discriminable, one might not notice the temporal order of these parts *as* temporal. The parts of a thought are connected by syntactic relations, and these relations must correspond to semantic relations between constituents of the content of which the subject is aware, if the subject is to be aware of the content as a single proposition. But then, what we might call the logical or semantic structure of the propositional parts—the way they are experienced as bound together in a unified whole—could mask their experienced temporal order. Provided the transition from one part to the next is fast enough, there may not be any introspectively noticeable “time lag” between, e.g., beginning to think about Ricky, and beginning to think some particular thing about Ricky, a gap in which one would have the opportunity to pause and wonder to oneself, so to speak, “what *about* Ricky?” (the way one might when reading a sentence about Ricky that began at the bottom of one page but finished at the top of the next). As a result, even if there *is* a split-second transition between, e.g., thinking of an object and predicating something of that object, the thinker’s awareness of the unity between the parts of her thought might be sufficient to obscure this fact. To the degree that the subject notices the temporal relations, she is inclined if pressed to interpret them as semantic relations, i.e. as certain parts of the content taking semantic

or logical *priority* over other parts, in the act of fitting the parts together. Thus, on reflection she seems to be aware only of a structured whole, rather than of a series of temporally separate parts.

If what I just described is on the right track, there could be cases in which neither first order nor higher order awareness affords the subject access to the proposition *the way she actually becomes aware of it*: in successive connected conceptual units. This does not imply that the subject fails to be aware of these conceptual units successively—she is aware of them successively, just not aware of them as a succession. We should conclude that the mere fact (if it is a fact) that it can seem to us that a thought arrives all at once is not sufficient to establish Simultaneity or Simultaneity*. But then we lack an adequate a posteriori reason to accept Premise 1 of the Process Argument.

V. Conclusion

I have surveyed two versions of the Process Argument and found that both fail to provide adequate support for the argument's first premise. The A Priori argument does not justify the supposition that a subject's entertaining a unified propositional content must happen in a single moment, and the A Posteriori argument rests on dubious assumptions about how our thoughts introspectively seem to us and what this can tell us about their actual temporal structure. Thus, we should conclude that PA fails to show that thoughts are not experiences. Unlike the Simple Response, this second response to Tye and Wright's argument succeeds in defending all versions of the Cognitive Phenomenology thesis from their attack.

This result means we are free to accept or reject the existence of cognitive phenomenology based on the evidence we will consider in the rest of the thesis. But in any event, our stance should not be based on the considerations presented in the Process Argument. In the remainder of the thesis, I will leave open the question of whether thoughts are best understood as states, events, or processes, and I will not worry much about which of these terms I use to speak about thoughts. In Chapter 3, finally, we turn to consider the main arguments in favor of a proprietary phenomenology of thought.

Chapter 3

Phenomenal Contrasts and Self-Knowledge

I. Introduction

In Chapter 2, we saw that there is no clear incompatibility between the temporal natures of thought and phenomenal experience that might pose an immediate obstacle to the possibility of a phenomenology of thought. Tye and Wright's argument for such an incompatibility proved to be specious. In this chapter we will survey the major positive arguments in the philosophical literature for accepting the existence of cognitive phenomenology.

The primary purpose of this chapter is to argue that, while the balance of the evidence does, just about, favor some version of the Qualitative Thesis, the sort of cognitive phenomenology for which the Liberal arguments provide evidence falls well short of supporting the Modal Thesis. Establishing this conclusion will set the stage for the second half of the thesis—Chapters 4 and 5—in which I first elaborate on what I take to be the most plausible version of CP, and then present my main argument against the claim that cognitive phenomenology is capable of determining the contents of thought. What we will be left with is a picture of the phenomenology of thought that at best only vindicates a Moderate CP position.

There are two central arguments in the literature in support of the existence of CP: the Phenomenal Contrast Argument and the Self-Knowledge Argument. In this chapter we will examine both in turn, seeking to answer (a) whether, and to what degree, the arguments are successful, and (b) if the arguments are successful to some degree, which versions of CP they support.

A secondary ambition of this chapter is to draw attention to the way these two arguments relate to each other in the case for CP typically presented by Phenomenal Liberals. The Phenomenal Contrast Argument offers a degree of evidence for some type of Moderate (Qualitative only) CP Thesis. But, as we will see, on its own it can offer little or no support for the Modal Thesis, and so must be supplemented with further argument if it is to be used to defend a Strong (Qualitative plus Modal) CP position. The Self-Knowledge Argument, on the other hand, presented on its own as a reason to accept CP, is in a much weaker dialectical position, since Conservatives can dismiss it outright by rejecting the view of introspection on which the argument explicitly depends. However, what I want to suggest is that the Self-Knowledge Argument—or the intuitive motivations driving the argument—taken as a *supplement* to the Phenomenal Contrast Argument, offers the chance to bridge the gap between the latter argument and a conclusion that actually supports the Modal Thesis—provided one accepts those intuitive motivations. And indeed, these motivations are precisely what many proponents of Strong CP are implicitly or explicitly relying on when they appeal to phenomenal contrasts to argue for both the Qualitative and Modal Theses, as I will try to show.

But this is a rather long story to tell, and it will help to consider several types of contrast arguments in detail before finally turning to the relationship between these

arguments and the appeal to self-knowledge. In the rest of the chapter, I will look at three versions of the Phenomenal Contrast Argument in sections II-IV, before turning to the Self-Knowledge Argument in section V.

II. Understanding Experience and Sensory Imagery

1. Examples of the Phenomenal Contrast Argument show up in an impressive variety of forms in the literature, but each instance shares a common strategy, viz. to point to (actual or hypothetical) cases in which there is arguably a phenomenal difference between the experiences of two subjects, or two experiences of the same subject, and then to argue that we need to appeal to *sui generis* cognitive phenomenology in order to give a satisfactory account of this “phenomenal contrast.”⁵⁵ The most common Conservative response to this strategy is to argue that any alleged instance of CP in these cases can be accounted for in terms of differences in non-cognitive, specifically sensory, phenomenal properties.

I will assess the state of the current debate over the standard Phenomenal Contrast Argument by focusing on the dialectic between Charles Siewert and Jesse Prinz. I think this focus is appropriate given that Siewert’s argument is one of the strongest in the literature, and Prinz’s reply is exemplary in the uncompromising and ingenious way it attempts to account for all of the relevant data in terms of sensory imagery. I take it these two philosophers’ arguments are sufficiently representative of Liberal and Conservative strategies generally.

⁵⁵ In addition to the arguments mentioned in this chapter, see Horgan & Tienson 2002; Horgan, Tienson, and Graham 2003; Kriegel 2015; Pitt 2004; Strawson 2011; and Nes 2012 for a representative sample.

Siewert's route to Liberalism proceeds by his first noting the existence of a contrast between what it's like to understand the meaning of a word or sentence and what it's like to hear or read the same word or sentence without understanding. For instance, a classic example of this sort of "understanding experience" phenomenal contrast is due to Galen Strawson (1994): there is surely a qualitative difference in what it is like for us to hear someone say something in a language we do not understand, say French, versus what it would be like to hear the same utterance if we understood the language (though such a case somewhat overcomplicates the issue, since the two experiences, of hearing the French sentence without and with understanding, cannot be experienced consecutively by the same person). Siewert makes a strong case for much the same point by asking us to consider the experience of reading a difficult passage, such as the following one from Thomas Jefferson:

Every difference of opinion is not a difference of principle. We have called by different names brethren of the same principle. We are all Republicans, we are all Federalists. If there be any among us who would wish to dissolve this Union or to change its republican form, let them stand undisturbed as monuments of the safety with which error of opinion may be tolerated where reason is left free to combat it. (2011, 251)

We can imagine getting to the end of the final sentence and realizing that we have not followed what we have just read, even though we are fluent English speakers and are familiar with all of the words. We start again with greater concentration, and this second time grasp the meaning of the passage. It seems hard to deny that (at least in a great many cases) what it is like the second time we read the sentence, with understanding, is different from what it is like the first time, though in each instance the experience of reading could be accompanied by the very same sensory phenomenology, e.g. the visual experience of the words on the page and the auditory imagery of inner speech. Nonetheless, the noted

difference in the two experiences is a phenomenal one. Thus, Siewert contends that this phenomenal difference must be due to the presence of distinctive *non-sensory* phenomenal character, which is tied to our comprehension (250-258).

Siewert's general strategy is to argue that his opponent must be able to point to some phenomenal character that (a) completely captures the difference in what it's like to be in the mental state we are in when we read the sentence the second time with understanding, and yet, since it is supposed to be non-cognitive, (b) *could* be enjoyed by a subject who utterly fails to grasp the sentence (253). But all the plausible candidates for this phenomenal character to which the Conservative might point are phenomenal properties that could be instantiated in a subject *before* comprehension in a possible contrast case scenario, and thus could remain invariant across the transition from non-comprehension to comprehension, failing to account for the phenomenal change. In this way, Siewert carefully addresses various possible descriptions of the contrast, and argues fairly persuasively that each is inadequate to explain the phenomenon (252-258).

For instance, he contends that the phenomenal contrast between failing and succeeding to understand language cannot be entirely accounted for by positing the addition of a "sense of fluency" in understanding: experiencing a sentence as grammatically well-formed, properly parsing it, reading it with fluidity or a certain intonation, and so on. While Siewert is aware of the phenomenal differences that can be due to such factors, he points out that we often fail to understand the meaning of a difficult passage while reading due to a slight lapse in concentration, even as we continue to read the passage to ourselves with the same parsing and prosody we would exhibit in the event of fully comprehending the passage. In such cases we often experience our inner voice reading *as if* we understand,

but we “draw a blank” and abruptly realize that the content is missing. But when we then go back and read to ourselves in exactly the same manner but with greater concentration, there is a novel phenomenal element that accompanies our grasping the content. So we are left with having to explain this phenomenal difference in terms of the difference itself between grasping and not grasping meaning or content.⁵⁶

He presents several other phenomenal contrast cases along the same lines. For instance, he points to the phenomenon of interpreting what someone has said one way, then suddenly becoming aware of a second interpretation, as when someone says to you, “I think I have a virus,” and you first think they are saying they feel sick, but suddenly interpret what they are saying as a statement about their computer (261). There is a phenomenal difference between the experiences of the two interpretations. Another example is the feeling accompanying a sudden realization of what precisely an interlocutor is referring to, for example if someone asks you, “do you have the book for me,” and it takes you a second or two to identify in thought the book they must mean (258). He thinks that in both of these types of cases the Conservative is going to have trouble, because variation in perceptual experience and sensory imagery alone cannot always account for the experience of a sudden shift in meaning or a sudden grasp of meaning. Thus, Siewert concludes we should reject Conservatism and embrace Liberalism.

2. In responding to Siewert, Prinz’s goal is to argue that all the phenomenal differences we have considered can be accounted for by differences in sensory phenomenology. For

⁵⁶ Note that the ability to deal with this point about a “sense of fluency,” and related objections, makes Siewert’s phenomenal contrast cases stronger than Strawson’s foreign language contrast case, since considerations of grammatical parsing and prosody arguably *can* explain the latter case.

instance, when I realize that by ‘virus’ my coworker could mean that his computer has a problem rather than that he is sick, a visual image of a computer flashes into my mind; when I remember what book my friend is referring to, I experience an image of the book’s cover; and when I am reading with full understanding, all sorts of sensory “simulations”—sensory images symbolizing the contents of my thoughts—float through my consciousness, while these images are absent when I simply recite the words to myself without understanding what they mean (Prinz 2011, 185).

One might wonder: what about the possibility that some thoughts are made up of concepts that are too abstract or complex, or are otherwise unsuitable, for representation by sensory imagery? And what if such thoughts flash through the mind at a rate that would make it unwieldy for the subject to experience simulations at the same time? Consider thoughts about mathematics or logic, or about an abstract problem in philosophy. Or take the last sentence of the passage from Jefferson quoted above: perhaps *some* small portion of its meaning can easily be represented by images, but surely representing *all* of its meaning imagistically is difficult, and the idea that we do so every time we consciously grasp the meaning is implausible.

Prinz has a ready response to this worry (185-186). He points out that the objection neglects to consider *verbal* imagery: the inner speech spoken to ourselves in consciousness, perhaps along with visual imagery of written words. Even extremely abstract and rapid thoughts can be consciously experienced through verbal imagery, which is a type of sensory phenomenology. Both verbal and nonverbal imagery can function as a kind of “shorthand” that allows for quite complex concepts and thoughts to register in consciousness, without one having to consciously experience all of their complexity. For

example, the word ‘theoretical’ in inner speech can serve as a “label” in consciousness for a concept that plays a complex conceptual role in one’s thinking, even though *understanding* the concept involves more than merely consciously *experiencing* the label (186). In this way, any concept can be experienced and distinguished in consciousness through a combination of verbal and nonverbal imagery. No appeal to uniquely cognitive phenomenal properties is needed.⁵⁷ As Prinz puts it,

If this picture is right, inner speech and non-verbal imagery are mutually supporting resources that can each pick up for the other’s limitations. Words can help us experience very abstract thoughts or thoughts that arise too quickly to simulate in imagination. They can also disambiguate images, allowing us to think Tweedledee, when a visual image appears equally just like Tweedledum. Non-verbal images help us to disambiguate words and simulate what our words represent. For most of us, both elements are constant components of conscious thought (189).

3. I take it that in the *majority* of cases Prinz succeeds in making it at least plausible that his Conservative story is a viable alternative to Siewert’s explanation of understanding experience phenomenal contrasts. Nevertheless, I think Prinz’s reply to Siewert is ultimately inadequate. The weakness lies in his “division of labor” account of the sensory basis of cognitive experience (189). He claims, “verbal imagery can explain every instance of conscious thought that cannot be accounted for by appeal to images of the contents of our thoughts” (185). Every phenomenal contrast case can be explained by a difference in verbal imagery, non-verbal sensory imagery, or both.

⁵⁷ Both parties to the debate assume that sensory imagery is not metaphysically sufficient for thought or understanding—i.e. they reject a traditional empiricist view that affirms the Modal Thesis but not the Qualitative Thesis. So the claim that all that is needed to explain the contrast cases is sensory imagery of one form or another amounts to the claim that the contrast cases fail to support both the Qualitative Thesis and the Modal Thesis. Prinz is a self-confessed “neo-empiricist,” in that he believes thoughts have sensory “vehicles”; but he does not think instantiating phenomenology of any sort is sufficient for thinking (2002; 2011, 175).

Prinz's picture invites the following response. He concedes that there are thoughts whose contents do not lend themselves to being represented by visual imagery or any other kind of non-verbal sensory "simulations,"⁵⁸ and therefore that such thoughts when entertained are phenomenally conscious only by virtue of verbal imagery: inner speech and perhaps visual images of written words. But then surely there are possible cases just like Siewert's examples, but in which the sentence that the subject suddenly understands expresses content that would ordinarily be consciously entertained via verbal imagery alone.

As Siewert points out, it would seem the experience of reading or listening to a sentence without grasping its meaning could include all the same verbal imagery that the experience of understanding the sentence includes.⁵⁹ But then how should we explain the phenomenal contrast in this case? As we have just seen, Prinz's *standard* way of dealing with a case of phenomenal difference but same verbal imagery is to appeal to non-verbal imagery. But we have stipulated that the case at issue is one in which expressing the thought in question typically would not involve *any* non-verbal imagery, and it would be ad hoc for Prinz now to insist that there *must* be some, just to explain the contrast. On the other hand, if the phenomenology of the understanding experience is exhausted by its verbal imagery, then we are at a loss to explain why it inevitably differs from the phenomenology of experiencing the words without understanding them. At a loss, that is, if

⁵⁸ Examples of such thoughts might help here: *This only is certain, that there is nothing certain* (Pliny the Elder). *Scandal is gossip made tedious by morality* (Oscar Wilde). And from Prinz's own paper: "These findings show that conceptualization can influence perception in dramatic ways: there are shifts in prototypicality, verbal labeling, generation of associated images, and allocation of attention to category-relevant features" (183).

⁵⁹ Of course it is likely that the subject's comprehending the sentence will eventually *lead* to different thoughts and therefore additional imagery.

we are determined to adopt Prinz's policy of denying non-sensory, *sui generis* cognitive phenomenology.

There are several ways we might attempt to resist this argument on the Conservative's behalf. We might hold that when a subject comes to comprehend a sentence's meaning after first not comprehending it, and the comprehension is accompanied merely by verbal imagery, there is not necessarily any phenomenal difference before and after the moment of comprehension. But if we are inclined to grant the phenomena of phenomenal contrasts in general, it seems a dubious move to suddenly deny that there is a phenomenal contrast in certain cases, due to how abstract or complex the thought's content is or how rapidly the thought is tokened in such cases—unless we think it is introspectively clear that phenomenal contrasts only ever amount to an addition of non-verbal sensory images. Personally, I am not confident I can imagine any examples in which what it's like to understand and fail to understand are identical. Instead, we might deny that there are any conscious thoughts lacking non-verbal sensory imagery after all. But this seems implausible, and anyway Prinz has conceded that there probably are such thoughts.

A more creative conjecture goes as follows: just because a certain proposition is most naturally entertained in thought accompanied by verbal imagery alone, it does not follow that when one *comprehends the meaning of a sentence* that expresses that proposition, the only imagery involved is verbal. Entertaining a certain propositional content of one's own accord is one thing; understanding that same propositional content as the meaning of a sentence with which one is presented, and which one must make a deliberate effort to comprehend, is another thing. Perhaps, when one suddenly grasps the

meaning of a sentence with effort after first not understanding it, non-verbal sensory images necessarily flow through one's consciousness.

But this response, while resourceful, is also implausible. Grasping the meaning of a sentence constitutively involves entertaining a thought whose content is that meaning.⁶⁰ It is a way of exercising of the faculty of thought. If we admit that a proposition can be consciously thought with only verbal imagery, it would be strange for us to insist that consciously comprehending a sentence that expresses that same proposition *must always* involve more than verbal imagery. I can see no principled reason for insisting this is necessary in the one case if it is not in the other.

Perhaps the most interesting response to the problem posed by these contrast cases is to claim that somehow the specific *kind* of verbal imagery that accompanies the grasping of linguistic meaning is unique to states of genuine comprehension and is not normally duplicated in sensory states in which one experiences words without understanding their meaning. Obviously much more would have to be said about what makes the verbal imagery of genuine cognitive states unique, and this would have to be some phenomenal feature or other. However, if their uniqueness were to come from some *non-sensory* phenomenal properties that somehow affect or interact with these sensory images, then this fact would seem to support (at least) a Moderate CP Thesis, and so would be no help to Prinz. In chapter 4, I will describe a kind of phenomenal feature along these lines that may be typically present in instances of genuine understanding and typically absent when a subject fails to grasp a thought.

⁶⁰ Siewert also argues that "occurrent understanding [of written or spoken speech] is a form of episodic conceptual thinking," on the grounds that understanding our *own* spontaneous (inner or outer) speech is a form of thinking, and understanding others' speech involves essentially the same mental capacities (2011, 248).

Perhaps these considerations will not move the committed Conservative. There are several points in the argument where the Conservative can dig in her heels if she chooses: e.g. she may claim that Prinz is too concessive to his opponents, or allow that the phenomenal somehow extends beyond the sensory imagery to which Prinz appeals, while nonetheless failing to include CP. So the standard Phenomenal Contrast Argument seems to me to be less than a knock down argument. But as things stand, once we adopt Prinz's strict commitment to explaining all phenomenal differences in sensory terms—along with his other assumptions—it is difficult to know what to say about the particular type of phenomenal contrast case that I have described, in which the conscious grasping of content does not involve any non-verbal imagery. So I think we should conclude these cases provide some evidence for the existence of a (perhaps merely generic) kind of non-sensory understanding phenomenology.

4. Before moving on, we should pause to consider which versions of the CP thesis could be supported by a standard, Siewert-style contrast case argument. Siewert himself does not distinguish between Modal and Qualitative Theses, and he seems to take his argument as supporting both versions of the CP view. He states his main target as the view that the experience of thinking is reducible to “what it's like to possess concurrent, separable sensory features—i.e., ones whose possession at that time is insufficient to guarantee anything about our thinking and understanding as we then do” (248). So he is endorsing both the claim that the phenomenal features in question are not sensory (hence the Qualitative Thesis) and the claim that they are not “separable” from actual thinking and

understanding (hence the Modal Thesis).⁶¹ But I think on the simplest understanding of the standard Phenomenal Contrast Argument, it only straightforwardly supports the Qualitative Thesis, since these contrast cases do not show that the phenomenology that comes with understanding necessitates thought; rather, our consideration of these cases only suggests that the contrasts cannot be explained by appeal to any change in sensory phenomenology.

One might think we can establish the Modal Thesis by consideration of these contrast cases as well, by a kind of inference to the best explanation. Siewert challenges his opponent to find a phenomenal duplicate of the experience of reading the passage with understanding, which nevertheless could occur separated from such understanding. So, if all the various sensory phenomenal properties to which we might appeal in order to meet this challenge apparently fail—because, while they *are* separable from understanding, they do not explain possible contrast cases in which those very features are present *before* the contrast—then we might suppose that the best explanation for this failure is that there does not exist a phenomenal feature that can meet Siewert’s challenge. The phenomenology of the understanding experience must be sufficient for the occurrence of that very episode of understanding (i.e. thought), which is just to say the Modal Thesis must be true.

But have we been given enough reason to think that this really *is* the best remaining explanation for the phenomenal contrasts? The problem is that it is very difficult to show, simply by considering the obvious examples of visual and auditory perceptual experience

⁶¹ Strictly speaking, this passage could be read as implying that any phenomenology that is *either* non-sensory *or* inseparable from thought should count as CP. But in context it is clear we should read Siewert as defending the position that there exists CP that has *both* these features.

and imagery, that there are no other plausible explanations involving phenomenology that is “separable” from the activity of understanding that we have not yet addressed. In particular, while we may be reasonably confident that we have ruled out any other sensory alternatives to explain the contrast—since we have a reasonably circumscribed understanding of what counts as sensory phenomenology—our argument has not ruled out possible *non-sensory* phenomenal features that are insufficient for, could be separated from, occurrent thought and understanding—features that vindicate the Qualitative, but not the Modal, Thesis. Thus, at least based on what we have said so far, it is unclear how Siewert could be at all confident that his challenge cannot be met.

Indeed, this issue for Siewert’s argument points to a quite general problem for any Liberal attempt to establish the Modal Thesis by appeal to contrast cases: how could a modal conclusion *ever* be justified simply by consideration of situations in which some phenomenal change accompanies the beginning of a thought? For how could claims about what is or would be the case in any number of imagined real world scenarios justify our reaching a blanket conclusion about what *must* be the case in any possible instance of conscious thinking?⁶² Reflection on these scenarios may tell us that they involve the addition of phenomenal character distinct from hitherto recognized phenomenal kinds, and perhaps may suggest some details about the qualitative nature of this character, but it is difficult to see how reflection on the introspective evidence alone could indicate to us what modal properties, vis-a-vis thought, the phenomenology in question exhibits.

This challenge is one that must be faced by any contrast case arguments we might consider, provided the argument is meant to support not just the existence of non-sensory

⁶² Thanks to Tim Bayne for eventually succeeding in impressing upon me the importance of seeing the issue in this way.

phenomenology that can account for the contrast, but also a necessary tie between that phenomenology and thought. I do not mean to deny the possibility of additional argumentation that could supplement the appeal to contrast cases and motivate acceptance of the Modal Thesis. But we will hold off on this issue until section V below.

III. Higher-level Perceptual Phenomenology

1. So far I have argued that the first main Liberal strategy for establishing the existence of cognitive phenomenology, the appeal to phenomenal contrast cases, gives us a modest, defeasible reason to accept a Moderate CP position, over the outright denial of CP. Next, I want to consider two variations on the Phenomenal Contrast Argument. In section IV we will look at contrast cases that supposedly involve a type of CP linked to the phenomenology of agency, what I will call *agentive-cognitive phenomenology*. And in the present section we will consider a contrast case argument that appeals to alleged top-down effects of concepts on the phenomenology of perceptual experiences.

I have two principal reasons for wanting to examine other contrast case arguments in addition to the standard, Siewert-style cases. The first reason has to do with the observation I wish to make about the relationship between contrast arguments and the argument from self-knowledge, which I will take up in section V. Since I want to make a general point about how these arguments are often or typically related, I need to consider more than just a single example of a Phenomenal Contrast Argument in action, so that we might observe a pattern emerging.

Secondly, while I do not pretend to be providing an exhaustive survey of the variations of this argument in the literature, I want to show that the conclusions that I drew

about the Phenomenal Contrast Argument while considering Siewert and Prinz's arguments are not a mere artifact of my choice to focus on those particular authors. Contrast cases like Siewert's that involve linguistic understanding or realization experiences may be more obviously relevant to the question of CP than the cases we will be considering below. But for anyone who is unsure about whether she really enjoys anything like the sorts of experiences that Siewert describes, contrast cases that focus on *other* sorts of experiences—perhaps less obviously linked to thought but with more vivid and less controversial phenomenology—may prove intuitively more gripping. So if a case can be made that these sorts of experiences involve cognitive phenomenology, these contrast cases could provide a kind of back door entrance to the CP position. And we ought to consider whether the things I said about Siewert's argument apply to these cases as well.

As I mentioned earlier, there is an ongoing debate in philosophy of mind and epistemology about whether conscious perceptual experience only represents “low-level” features of our environment apprehended by our senses—properties such as colors, shapes, and motion—or whether, in addition to these, perceptual experience can represent more complex, “higher-level” properties such as causation, affordances, or kind properties.⁶³ The view that says only lower-level properties are represented in perceptual experience, we can call the “sparse” view, while the view that countenances higher-level perceptual properties as well can be called the “rich” view. The question of whether perceptual experiences can represent higher-level properties, i.e. whether they can have “rich” content, and the question of whether some perceptual experiences involve cognitive

⁶³ Sounds, smells, temperature, and so on, are picked up by other modalities. In what follows I will mostly restrict myself to talking about visual experience.

phenomenology over and above their sensory phenomenology have not always been clearly distinguished.

What is it for perceptual experience to represent a given property? First, perceptual experience has phenomenal character. So the fact (if it is a fact) that empirical psychology suggests that perception represents higher-level kinds and categories, is not yet enough to show that perceptual *experience*—that part of the deliverances of perception that constitutes phenomenal consciousness—has such content. Second, we can say that if “something can be seen [heard, felt, etc] to have” a given property, then that property can be represented in visual (aural, tactile, etc) experience (Siegel 2006a, 481). If we go on to accept both this and its converse (as seems reasonable), then it follows that the properties represented in perceptual experience, those that make up its content, are all and only those that—strictly and literally—are *perceived* (seen, heard, felt, etc.). This characterization might seem trivial, but it helps us distinguish visually represented properties from properties that we merely judge to be instantiated *on the basis of* what is visually represented. For example, I might say that I can see my friend is angry with me for a thoughtless comment, but what I mean in this case is that I *judge* that this is so, based on the evidence of what I do in fact see, e.g. her furrowed brow and frown.

Proponents of the view that perceptual experience is rich often appeal to phenomenal contrast cases. For instance, arguably there is a phenomenal difference between the experiences of spotting a bird of a particular species in a nearby tree, or hearing its distinctive birdsong, as a first-time bird watcher or an avid bird watcher. And there is a clear phenomenal difference between the experiences of looking at the famous duck-rabbit drawing and seeing a duck or looking at it and seeing a rabbit. I will call these

sorts of phenomenal contrast cases perceptual-recognitional (PR) phenomenal contrast cases, since they involve a difference in what the subjects recognize based on their perceptual experiences.

Here is an especially striking example from Susanna Siegel (2006a):

Suppose you have never seen a pine tree before, and are hired to cut down all the pine trees in a grove containing trees of many different sorts. Someone points out to you which trees are pine trees. Some weeks pass, and your disposition to distinguish the pine trees from the others improves. Eventually you can spot the pine trees immediately. They become visually salient to you. Like the recognitional disposition you gain, the salience of the trees emerges gradually. Gaining this recognitional disposition is reflected in a phenomenological difference between the visual experiences you had before and after the recognitional disposition was fully developed. (491)

Siegel argues that the phenomenal difference that accompanies the gaining of recognitional dispositions supports the claim that visual experience has rich representational content, an addition of which explains the phenomenal difference in question.

What is the relationship between this kind of argument and a phenomenal contrast argument for CP that uses similar examples? In chapter 1, I proposed we understand concepts to be all and only those representations that can serve as the constituents of thoughts, and thoughts as those mental particulars produced by an exercise of the faculty of thought (which I went on to describe in greater detail). In order for a Siegel-style PR contrast case argument to support CP, the argument would have to show that the experiential contrast in question is essentially conceptual—i.e. constitutively involves differences in representational types, tokens of which serve as constituents of thoughts—*and further*, that the deployment of concepts within this experience essentially involves the operation of the faculty of thought. If a perceptual experience with conceptual content

were to necessitate the thinking of a thought in this way, then the experience would vindicate the Modal CP Thesis.

One might think that if the perceptual experience that accounts for a given phenomenal contrast is essentially conceptual, it is also likely to require the exercise of thought. In what follows I'm going to make this simplifying assumption. I'm also going to assume that if a perceptual experience essentially involves the representation of higher-level properties, this fact provides good reason to suppose the experience is essentially conceptual. Neither assumption is unproblematic, but in granting them here I'm only strengthening the position against which I will be arguing.

What we are left with, once we grant these two assumptions, is a situation in which, *if* the PR phenomenal contrast cases succeed in supporting the rich view of the content of perceptual experience, they thereby support the existence of cognitive phenomenology as well. For a perceptual experience to represent high-level properties would require that at least some of the phenomenology of that perceptual state be cognitive phenomenology, in the sense of the Modal Thesis. And since this phenomenology arguably would be non-sensory (in the sense described in chapter 1), it could also count as CP in the sense of the Qualitative Thesis. Thus, the question of whether Siegel's PR phenomenal contrast argument succeeds becomes directly relevant to evaluating the CP debate.

2. The argument below is a variation of Siegel's argument for her higher-level content view (2006a), adapted to apply directly to the question of cognitive phenomenology. We start with a description of two experiences that together constitute a PR phenomenal contrast case, such as her pine tree example, where the difference between the "Before"

and “After” experiences seems to arise due to the subject’s acquiring certain recognitional abilities. We then reason as follows:

- P1** In a PR phenomenal contrast case, the subject’s overall After experiences differ phenomenally from her overall Before experiences
- P2** If the subject’s overall Before and After experiences differ phenomenally, then her *perceptual* Before and After experiences—which are parts of the overall Before and After experiences—must differ phenomenally
- P3** The phenomenal properties of perceptual Before and After experiences supervene (with metaphysical necessity) on their intentional properties: if there is a phenomenological difference between these two experiences, then they must differ in intentional content
- P4** If there is a difference in content between the two perceptual experiences, then this difference is a difference in the higher-level properties that each represents
- C1** At least some perceptual experiences—i.e. those Before & After pairs that constitute PR contrast cases—possess (and differ with respect to) rich representational content, i.e. they represent higher-level properties
- P5** In order for a subject’s perceptual experiences to represent higher-level properties, the subject must possess and employ certain concepts, which in turn requires the exercise of thought
- P6** Any two perceptual experiences of the same phenomenal types as a given pair of Before & After experiences will also represent higher-level properties

From C1, P5, and P6:

- C2** At least some phenomenal types are such that any perceptual experience that instantiates one of those types necessitates that the subject undergoing it exercises her faculty of thought

C2 entails the Modal Thesis: the instantiation of certain phenomenal properties

necessitates thought.⁶⁴

⁶⁴ See Carruthers 2000 for similar arguments, and critical discussion of these in Carruthers & Veillet 2011.

3. Let us evaluate the PR Contrast Argument by looking at each of its premises. P5 just articulates our two assumptions mentioned above, which allow us to move from a conclusion about higher-level content to a conclusion about thought. Since C1 is stated in terms of perceptual experiences, but the Modal Thesis is in terms of phenomenal properties, P6 is needed as a bridge between talk of experiences and talk of phenomenal properties. P6 is justified by the fact that the argument applies to experiences typed phenomenally: we can run the argument with any two experiences sharing the same phenomenology as the Before and After pair of experiences in our initial thought experiment. This leaves us with Siegel's four original premises: P1 through P4.

P1 just states the minimal assumption that there is indeed a phenomenal contrast in these cases, which seems hard to deny. One might worry that Premise 2 undermines the Liberal's attempt to co-opt Siegel's argument, since P2 can be taken to suggest that the phenomenal contrast is due to a difference in perceptual phenomenology *instead of sui generis* cognitive phenomenology. Indeed, in defending P2, Siegel focuses on refuting the idea that the phenomenal contrast is due to a change in CP. However, we have granted for argument's sake that any representation of higher-level content in perceptual experience should count as an instance of conceptual thought. Admittedly, this assumption has the slightly odd consequence that certain experiences we are classifying as *perceptual* may turn out to be necessarily dependent on the exercise of thought, and therefore include *cognitive* phenomenology. But I see nothing deeply problematic about this result.

Instead, one could resist the PR Contrast Argument by rejecting the third premise, that the phenomenal contrast reveals a difference in content between the two experiences. This would be to deny the popular doctrine that every difference in phenomenology

necessarily entails a difference in representation, and to allow that, in experiences of perceptual recognition at least, there is “mental paint”: phenomenal character that makes no representational difference (Block 1990). Since the majority of both Liberals and Conservatives accept some version of the claim that phenomenology supervenes on intentional content, I will not object to P3 here, but I return to consider it below.

This brings us to Premise 4, which says the difference in content must involve the representation of higher-level properties. Siegel’s defense of this move fails to explicitly engage with one of the most plausible accounts of the phenomenal difference, namely an appeal to shifts in the patterns of attention in the visual experience of the forest before and after one learns to distinguish pines from other trees. Others have already made this case so I will be brief.⁶⁵ The idea is that what changes when one employs new recognitional capacities are the features of our environment on which we visually focus and to which we devote greater cognitive resources. Although these changes are often instigated unconsciously and automatically, they can make a significant difference to how a scene—composed of arrangements of sparse properties like shapes and colors—looks to the subject. For instance, when we come to recognize pine trees, we might visually attend to the distinctive shape or coloring of their tree bark, and when we see the drawing as a duck, we focus on the part of the drawing that corresponds to a duck’s face and bill.

Differences in experience brought about by attention are not likely to be sufficient for there to be differences in higher-level content, nor require that a subject acquires any new concepts. To be sure, differences in the exercise of conceptual capacities may be what *cause* the phenomenal changes, since “when one acquires a new concept one learns to

⁶⁵ See, for instance, Carruthers & Veillet (2011) and Levine (2011), who draw on the work of Pylyshyn (2003).

attend preferentially to those features of its instances that are distinctive of them” (Carruthers & Veillet 2011, 39). But this account offers a way for the use of concepts to have an effect on visual experience—by affecting visual inputs and how visual information gets processed—without necessarily being a *constitutive* part of visual experience itself. Thus, the “Attention Story” is able to offer an empirically plausible Conservative explanation as to why we should expect a phenomenal contrast in these cases.⁶⁶

Liberal defenders of the PR Contrast Argument can respond to the Attention Story in a couple of ways, but the best response, I think, is to insist that even if these cases exhibit differences in attention, they probably exhibit other phenomenal differences *as well*, since the contrast between the experiences seems to involve more than just an attention shift.⁶⁷ This extra phenomenal difference is then supposed to be what determines the difference in higher-level properties represented. Galen Strawson insists on this point, claiming (about the experience of looking at the duck-rabbit):

Many judge that nothing changes ‘purely visually’ as the figure changes aspect from duck to rabbit, but we can allow that this isn’t so, that there is some such change, while insisting that it will be trivial, indeed probably undetectable, considered apart from the cognitive-experiential change from seeing the figure as a rabbit to seeing it as a duck. There may also be other changes of sense/feeling experiential tone as a result of the visual change...when the switch occurs. So be it. None of these will add up to the experiential change that interests us, the change that can’t be adequately characterized without some reference to experienced duckishness and rabbitishness. (2011, 302)

What the Attention Story might seem to leave out is the feeling that I am recognizing something that (a) I *already* know about and (b) is *in* the visual scene itself—I am *seeing it*

⁶⁶ It is an open question exactly what one should say about the representational differences that such shifts in attention bring, or indeed whether the phenomenal differences brought about by attention correspond to representational differences at all (if not, attention phenomenology would be a kind of mental paint, and the Attention Story could be offered as an objection to P3 instead).

⁶⁷ Cf. Bayne 2009. Bayne offers several other responses to the Attention Story, thought space constraints prevent me from addressing them here.

there. One might think the shift in attention is only what is required to get the picture to “look right,” *in order* that then we can experience it as looking like something with which we are already familiar, as “matching” a stored representation. And it is this latter sort of experience that the Attention Story seems to have trouble explaining: the experience is not merely of the visual scene changing, but of it changing into something we recognize. And the recognition is arguably not just a “raw” rush of familiarity feeling on top of the visual experience, but is, so to speak, built right into the visual experience, since it involves the thing recognized being *perceived* in the scene itself. It is hard to imagine how a person whose seeing involved all the right patterns of attention but who utterly lacked any past exposure to rabbits could experience this sense that a rabbit is seen *in* the picture, or (to use a different preposition) that the picture is seen *as* a rabbit.

I do not pretend that this consideration will convince many who hold to the sparse view of perceptual content, but framing the point this way helps pinpoint why someone might find the Attention Story phenomenally inadequate. However, *even if* shifts in patterns of visual attention cannot capture all of the phenomenal differences in these cases, we should not necessarily conclude that the phenomenology at issue represents higher-level properties or involves the exercise of concepts, as I now will argue.

4. In defending P4, Siegel seeks to cast doubt on alternative accounts of PR phenomenal contrasts. The main alternative she considers is the idea that what gets added to the representational content of one’s visual experiences when one learns to recognize, e.g., pine trees, is what she calls a “pine-tree-shape gestalt” (499). A pine-tree-shape gestalt is “a complex of shapes—leaf shapes, trunk shapes, branch shapes, and overall pine tree

shapes... [it] is general enough that it can be shared by differently looking pine trees. But it is specific enough to capture the look shared by exemplary pine trees. The pine-tree-shape gestalt is invariant across differences in shape of particular pine trees.” The idea is that after one becomes an expert tree spotter, one’s After experience represents this complex shape property, rather than the content *pine tree*, and this addition to the content of one’s visual experiences is accompanied by a phenomenal addition that creates the phenomenal contrast. Representing the pine-tree-shape gestalt plausibly does not require having a concept of pine trees, nor does it require having a concept of the shape as a shape that a certain set of trees all have in common, but rather “it is enough simply to represent the respects in which various pine trees are in fact similar” (499).

It seems to me we could fill out the details of this somewhat vague proposal in a number of different ways.⁶⁸ Still, I want to say a few things in defense of pine-tree-shape gestalts as a way of objecting to P4. Since the account could involve the representation of any complex of sparse properties that are invariant across the kind of objects that are recognized in PR experiences, we can call this account of PR contrast cases the “Complex-Sparse-Gestalt Story”, or just the “CSG Story” for short.

Siegel points out that acts of perceptual recognition involving higher-level properties have a certain structure: roughly, the subject must have some sort of perceptual representation stored in memory, along with a sensory input representation, and the subject’s cognitive system must register a “match” between the input and the memory (498-499). This seems right. She goes on to say, “part of what’s at issue in the debate about [higher-level properties] is whether visual experience is an input to such processes

⁶⁸ For instance, I am not sure exactly what it means to call this complex shape property (or “complex of shapes”) a *gestalt*. See Chudnoff 2013 for a thought-provoking discussion of CP and gestalt experiences.

of recognition, or an output of such processes” (499). I take it what she means is that if we have reason to think this “matching” process is a constitutive part of perceptual experience itself, rather than something that happens afterwards, then this supports her view that such experiences have higher-level content. For if the matching is part of such an experience, then it makes sense that the experience would represent the perceived object as being of the kind, having the higher-level properties, of whatever is represented by the memory representation.

However, this picture of the structure of recognition fits just as well, if not better, with the CSG Story. CSG properties are just the sorts of properties we might expect a stored representation based on previous perceptual experiences to represent. So even if the recognition is part of visual experience in these cases, this does not point towards the experience representing higher-level properties *as opposed to* CSG properties, at least not without further argument.⁶⁹

For instance, the experience of the aspectual shift from duck to rabbit could be due to a shift in the visual memory representation (from duck image to rabbit image) that is matched to the visual input representation. I think this adequately explains the sense we have that we see a rabbit “in the picture.” One experiences the picture as having the same complex shape property as an image one has stored in memory, which is just to say one experiences the picture as resembling something one has seen before.

⁶⁹ Presumably this further argument would have to include reasons to think that the memory representations that get matched to perceptual inputs contribute higher-level content to an experience’s overall content. To try to meet this challenge one might adopt concept empiricism, according to which concepts, such as PINE TREE, employ perceptual-based representations as vehicles (Barsalou 1999, Prinz 2002). But in fact, concept empiricism is consistent with there being representations stored in memory that are identical to those representations that serve as conceptual vehicles with respect to their lower-level content and phenomenal character, but which lack the higher-level content of conceptual representations (Prinz 2002, 2011).

Crucially, the stored perceptual representation would not have to be a concept itself, but could be exclusive to perceptual systems. And, one would not even need to have a concept (or indeed any other representation) of ducks or rabbits in order to be able to have such an experience, since what we recognize “in the picture” seems to be a certain superficial *look* or *appearance* with which we are familiar, rather than anything so complex and abstract as a natural kind property (cf. Prinz 2011).⁷⁰ Arguably, we can imagine a subject having been hooked up in such a way as to have contrast experiences while looking at duck-rabbits without actually being able to represent *rabbit*, provided she has memories of other things that had the same rabbit-like “look”. While application of kind concepts might still be linked *causally* to the process of matching memory representations to perceptual inputs, these considerations show it is at least questionable whether representation of kinds is constitutive of the PR phenomenal contrast experiences.

Siegel presents an objection to the Complex-Sparse-Gestalt Story, but, as she acknowledges, it is hardly conclusive. She argues that even if the CSG account were to work in the case of learning to recognize pine trees, it would not be of general use in all such cases, for instance the case of learning to recognize a frequent facial expression of an individual or group. Suppose the facial expression expresses doubt. One *initially* could learn to recognize the expression, say a certain strange contortion, without knowing it expressed doubt. This is the sort of recognitional ability that might be accompanied by one’s visual experience coming to represent a CSG property of the facial contortion. But,

⁷⁰ Thus, when Strawson asserts that the phenomenal shift in the duck-rabbit case “can’t be adequately characterized without some reference to experienced duckishness and rabbitishness,” the defender of the CSG Story can agree, provided by this description we have in mind superficial properties like looking-like-such-and-such (i.e. like a duck) and looking-such-and-such (i.e. like a rabbit), or having-such-and-such-shape (i.e. a duck’s shape) and having-such-and-such-shape (i.e. a rabbit’s shape).

she thinks that the additional ability to recognize the expression as one of doubt seems to be a change “that could be accompanied by a phenomenological change as well,” even though “it seems implausible to suppose that there must be a change in which color and shape properties are represented before and after one learns it is doubt that the face so contorted expresses” (499-500).

However, it seems to me likely that in such a case we would expect a subtle difference in the *emotional* phenomenology of our overall experience as we respond to seeing the facial expression having learned it expresses doubt (and presumably emotion experiences need not represent higher-level perceptual content). Siegel replies that we cannot rely on a shift in emotional phenomenal character here, since the person one observes with this look on her face “might not be anyone significant” for one, e.g. just “a talking head on television” (500). But I do not see why a shift in our emotional (and/or bodily-sensory) phenomenology when we “see doubt” in someone’s face should be dependent on them being someone we care about. It could be the emotional response is a natural result of our unconscious mindreading and empathic abilities.

Furthermore, the phenomenal contrast might correspond to some additional CSG content. For one might have stored representations of CSG properties that correspond to very general expressions of doubt one has encountered before, but which encompass more than just the property that one had previously recognized as characterizing this particular doubtful expression before one realized it expressed doubt. In other words, one might come to *see in it doubt*, when before one had only seen in it the same mysterious expression the person in question had made before, and this could correspond to one’s experience

coming to represent an additional CSG property. So Siegel has failed to mount a convincing reply to the Complex-Sparse-Gestalt Story.

Let us sum up. I have been at pains to show that there is no easy argument from the existence of Perceptual-Recognitional (PR) phenomenal contrast cases to the Modal Thesis, *even if* we grant that the presence of higher-level content in perceptual experience entails that that experience necessarily involves conceptual thought (Premise 5 in the argument above). I have attempted to refute the argument from the contrast cases to the Modal Thesis by attacking P4, the premise stating that the content that is added to the After experience in a PR contrast case must involve higher-level properties. To do so, I have outlined and compared three broad accounts of the phenomenal and representational change that takes place in these cases.

First, we could try to account for the phenomenal contrast in recognition experiences entirely in terms of changes in the subject's patterns of attention (the Attention Story). Second, we could posit, in addition to shifts in attention, an addition of complex sparse properties that are explicitly represented by the experience only after the subject recognizes something in his visual perception as being of a particular kind (the Complex Sparse Gestalt Story). Finally, we could account for the phenomenal difference between the Before and After experiences by supposing that the After experience involves an addition in the higher-level properties it represents. I have argued that the first and second accounts are at least as plausible as the third account. And since neither the Attention Story nor the CSG Story entails that the experience necessarily represents higher-level properties, the route from PR contrast cases to Strong CP is blocked.

5. Once again, as with the standard Phenomenal Contrast Argument we discussed earlier, it is worth considering which versions of CP the PR phenomenal contrast case argument is able to support. Unlike Siewert's argument, this argument (if it could be made to work) promises to support not just the Qualitative Thesis but the Modal Thesis as well. However, we are left to wonder how this support is possible, given the general challenge for any Liberal who seeks to use contrast cases to argue for the Modal Thesis, viz. the challenge of explaining how the mere fact of a phenomenal difference in these cases could support the conclusion that the phenomenal properties that account for this difference bring thought along with them in every *possible* situation in which they are instantiated.

I have given my opponent a leg up, so to speak, by granting that the representation of higher-level properties in experience is (metaphysically) sufficient for the use of concepts, and that concepts are likewise sufficient for thought (P5 in the argument above). We could reject these assumptions and the argument for the Modal Thesis would not go through. I think they are *prima facie* attractive claims, for which one could offer independent arguments. But the important point is that the PR contrast cases—and Siegel's initial argument for the conclusion that they exhibit a difference in the higher-level content of perceptual experience—must be *supplemented* with these additional premises that I have simply *assumed*, in order for the argument to have any chance of establishing the Modal Thesis.

Furthermore, even setting these additional assumptions aside, we are still left with a puzzle: how can consideration of a handful of hypothetical contrast cases show that the experiences in question are metaphysically sufficient for the representation of higher-level properties? The mere presence of the contrasts, on its own, cannot establish this claim.

The key first step can be traced back to P3, the claim that a phenomenal difference in the contrast cases necessitates a difference in the experience's content (or, as seems plausible in these cases, an *addition* in phenomenology necessitates an addition in content). This claim seemed innocuous enough when we considered it earlier, since I pointed out most of the philosophers on both sides of the debate will accept that phenomenal character supervenes on content, and will reject mental paint.

But now we can see that accepting this premise allows one to bridge the gap from consideration of phenomenal contrasts to a conclusion about the modal properties of the phenomenology that accounts for these contrasts. For if the mere presence of some novel phenomenology in these cases determines that the experience represents some novel content, then any account of what this content is likely to consist in is *ipso facto* an account of the likely modal properties of this phenomenology. But using the contrast cases to reason in this way depends on our prior acceptance of the general principle that there is no such thing as mental paint, and despite the widespread popularity of this claim, the phenomenal contrasts themselves do not establish its truth.

The moral of the story, once again, is that the proponent of phenomenal contrast cases faces a serious challenge to articulate how the phenomena of contrasts can offer us any justification for accepting the Modal Thesis. This challenge may not be insurmountable, but it will require significant supplementary argumentation, beyond simply pointing out the contrast and arguing that it cannot be accounted for in terms of sensory phenomenology with low-level content.

Do the PR phenomenal contrasts provide evidence for the Qualitative Thesis instead? Recall that in order for a phenomenal property to count as CP in the Qualitative

sense it must be non-sensory and (somehow) specially associated with thought. In chapter 1 we said that sensory phenomenal properties are all and only those perceptual phenomenal properties that are phenomenally *basic*, where this means the property constitutes an experience the having of which does not entail having any other distinct experience. If the Attention Story is correct, then I think the phenomenology in question is basic, since arguably the phenomenology of attention is partly constitutive of visual experiences of low-level properties, the having of which does not entail having any other distinct experiences (e.g. there is not some more basic visual experience that consists only of those same color and shape phenomenal properties, completely devoid of attention phenomenology). But if the CSG or higher-level property story is the correct interpretation of the phenomenal contrasts, then arguably the phenomenology that explains the contrasts is non-basic and thus non-sensory. Whether it should count as CP in the Qualitative sense would then depend on whether it is specially associated with thought or not. The proponent of this argument would have to make a case for this conclusion, but it is not crazy to think that this could be done. So PR phenomenal contrasts could provide us with a degree of evidence for the truth of the Qualitative Thesis.

IV. Agentive-Cognitive Phenomenology

1. Do we arrive at a different conclusion about CP if we appeal to agentive-cognitive phenomenal contrast cases instead of the sorts of cases we have focused on so far? Terry Horgan (2011) thinks that consideration of several “partial-zombie scenarios” in which subjects lack various kinds of agentive phenomenology will help philosophers skeptical or unsure about cognitive phenomenology to recognize that CP makes as much of a

phenomenal difference as agentive phenomenology does. By “agentive phenomenology” he means to encompass all those properties that make up the distinctive “what-it’s-likeness of experiencing one’s behaviors *as one’s own actions*” (62). By “zombies” he means imagined physical and functional duplicates of normal human subjects that lack phenomenally conscious experiences.⁷¹ But his creatures are “partial zombies,” meaning they only lack certain agentive or cognitive phenomenal features, and otherwise enjoy experiences with the full range of phenomenal character possessed by normal human experiences.

We are to imagine the experiences of these various partial zombies and compare them with the experiences of their normal human duplicates, as each are placed in various everyday situations. The expected result is that we will realize that the zombies’ experiences really are phenomenally impoverished in comparison to us in just the ways that Horgan describes, and that therefore the alleged phenomenal properties that they are stipulated to lack really do exist in normal experience. Clearly, this thought experiment is a variation on the method of using real world phenomenal contrasts to argue for the reality of CP.

Horgan invites us to consider several types of agency experience that have an unmistakably cognitive element. There’s *language understanding agentive phenomenology*, the phenomenology of experiencing the verbal behavior of oneself and others as speech acts, and one’s behavior caused by such verbal behavior as responses to linguistic communication. There is also *ulterior purpose agentive phenomenology*, the phenomenology of acting for purposes that go beyond simply the moving of one’s body in various ways, that is the experience of carrying out actions motivated by “occurrent desires

⁷¹ Horgan’s argument does not depend on his zombies being possible, but only conceivable.

that constitute ulterior motives,” and “occurrent beliefs about how to satisfy those desires behaviorally” (74). For instance, in a normal situation I might experience a desire to call my wife, and a belief that picking up my phone is a necessary step to fulfilling this desire, and therefore a further desire that I pick up my phone, and a further belief that moving my arm and hand to reach into my pocket is a necessary step to fulfilling this desire, and finally the action of moving my arm and hand *in order to fulfill the purpose of taking my phone out of my pocket and calling my wife*.⁷² By contrast, in the exact same situation, a zombie who lacked such phenomenology at most would experience only an intention to move his arm and hand in such-and-such a way, followed by this action.⁷³

We can classify all such alleged phenomenal properties as *agentive-cognitive phenomenology*. It is cognitive because it constitutes the experience of acting for purposes, which in turn entails that one consciously grasps these purposes (and, perhaps, the beliefs and desires that lead to the purposes). If we agree that a partial zombie who lacks some or all agentive-cognitive phenomenology is conceivable, and that there is an introspectively evident phenomenal difference between the zombie and ourselves, then it follows there exists agentive-cognitive phenomenology. In which case, there exists cognitive phenomenology.

2. What are we to make of this argument? In contrast to a “classic” phenomenal contrast case, such as the ones Siewert presents, Horgan’s partial zombie scenarios do not

⁷² This description is not meant to imply that in normal cases each of the occurrent beliefs and desires (more accurately: thoughts) that make up the chain of reasoning will be phenomenally conscious, only that some will be.

⁷³ Horgan also describes an even more impoverished zombie, who lacks even the experience of the intention to move his body in various ways, but I am simplifying here to focus on the lack of *ulterior* purpose phenomenology.

provide for us a contrast between two actual phenomenal states, which we might have experienced in real life. Rather, in imagining these zombies we have to *assume* that they are distinct from real experiencing human subjects, i.e. assume that there is some phenomenal properties over and above the sensory that the zombies are missing.⁷⁴ So Horgan's cases, by comparison, seem more contentious and less persuasive from the outset.

However, in spite of this disadvantage, at times I find these cases gripping in a way the classic contrast cases are not. Consider: my wife turns to me and interrupts my philosophical reveries by asking, "we are out of milk; can you go get some?" I say, "sure," get up from my chair, walk to the wardrobe to get my winter coat, pick up my wallet, put on my shoes, and walk to the corner store. I experience my wife's speech as a request for me to perform a particular action, to satisfy a particular desire of hers, which is further connected to other desires and intentions she has, which are at least somewhat transparent to me. Further, I do not experience each subsequent action as merely an isolated instance of my own physical behavior (intended or otherwise), but as steps in a larger process consciously directed at responding to her speech act by fulfilling the goal of acquiring a carton of milk.

The crucial question is whether the relevant sense in which I experience my behavior here as purpose-directed action, the relevant sense in which it "seems this way to me" is a *phenomenal* seeming. Initially, the fact that we can have (or imagine having) *sensory* experiences of our physical behavior without understanding speech directed at us,

⁷⁴ Horgan could reply that he is not asking us to assume the phenomenal difference, but merely betting that we will notice a difference between our own experience and the experience of the "zombie" subject that he describes for us. But then the success of the argument does not really depend *merely on the conceivability* of the "zombie" subjects, as Horgan seems to suggest, since they may be conceivable and yet phenomenally identical to normal subjects (this would explain their conceivability!). The success of the argument depends rather on our noticing that the imagined subjects are not phenomenally just like us, a task which is made difficult by the fact that normal humans do not experience this kind of contrast scenario in real life.

or grasping reasons for why we are behaving the way we are, suggests that there must be some non-sensory phenomenal extra that I have in this case but my zombie would lack. The difference between the experiences of purposeful actions and senseless behavior would seem to dramatize the fact that *there is something it is like to hold thoughts consciously before one's mind*, e.g., as purposes directing one's behavior. At least I can find myself having this intuition when considering Horgan's cases. What is the source of this intuition?

What Horgan himself says about such cases is revealing: whereas I experience a desire to comply with my wife's request and thereby form intentions to perform the various intermediate actions that will lead to my fulfilling that desire, my partial zombie's corresponding functional state "is experienced as the onset of a spontaneous, seemingly 'out of the blue,' desire" to move his limbs in various ways (2011, 71). This is supposed to follow from his experiencing my wife's utterance as "meaningless-seeming vocal noises emitted from her" (71).

We might challenge the assumption that this is how the zombie will experience her speech by pointing to the phenomena discussed earlier in the chapter, that the experience of hearing speech as meaningful language (quite apart from grasping its content) differs phenomenally from hearing the same speech as mere sounds. Given that the zombie possesses all the right functional states, we can assume that he will have the sensory phenomenology that corresponds to hearing her utterance as meaningful speech. Horgan might respond: the phenomenal effects of parsing spoken language, while perhaps sufficient to distinguish the zombie's experience from that of hearing mere noise, cannot by

itself explain why we hear the difference between parsed and unparsed speech as a difference between inherently meaningful and meaningless sounds.

However, one might still reasonably suppose that this awareness I have, that my actions are responses to what my wife has said, is explicable by the facts that (a) I experience rich sensory imagery directly after hearing her request (e.g. a visual image of a carton of milk as I hear the word ‘milk’), and (b) my auditory experience of hearing her request is directly causally connected to a non-conscious intention to walk to the store, along with subsidiary intentions accompanied by further imagery. If so, the fact that my zombie shares both (a) and (b) with me entails that he will not experience subsequent desires—e.g. to walk to the wardrobe and put on his coat—as coming “out of the blue.” Indeed, since the zombie is a functional duplicate of me, he *cannot* be any more puzzled by the intention to go to the store than I am! Furthermore, if we still worry that all the mental processing that connects his auditory experience to his intentions, and his intentions to his behavior, would occur strangely “under the surface” in the zombie in a way that it does not for me, that’s where the stream of sensory imagery comes in: imagery bridges the gaps in consciousness between an initial perceptual experience and subsequent experiences of our bodily responses to the perceptual stimulus.

In response, the defender of the zombie contrast cases might admit that the zombie’s being a functional duplicate ensures that he will not be caught off guard by his own behavior, but argue that this fact only explains why, in Horgan’s words, the zombie is especially “un-curious,” failing to wonder why he has the desires and intentions he has after hearing speech (2011, 71). It does not follow that the zombie’s intentional states, and his subsequent actions, do *not* seem to come “out of the blue” to him. As for the zombie’s

consciousness being filled by sensory imagery, one could argue that such imagery would not be enough to ensure that his inner cognitive processes are transparent to him the way mine are to me, rather than being opaque, since in the absence of agentive-cognitive phenomenology there would still remain an inner silence between one image and the next. The sort of awareness of my own motives that I enjoy could only be available to the zombie via additional *inferences* he must make to explain to himself the otherwise mysterious connections between the various images he experiences and his own behavior.

It should be evident that the disagreement just described in fact hinges on the role of phenomenal consciousness in self-knowledge. Horgan's opponent points out that one should only think that my partial zombie's intentions and actions come "out of the blue" if one thinks that a conscious subject—like a homunculus—only learns what is going on in her own mind by studying her phenomenology. But, so the opponent contends, a subject *isn't* only able to find out about her own mind this way. She gets to find out *directly*—because she *is* her own mind, so to speak, rather than a little man inside her mind. She merely needs to be access conscious, to be functionally set up in the right sort of way, in order to token the right sort of higher-order thoughts about her mental states and thereby be aware of them. So the view that Horgan seems to be relying on in order to get his contrast cases to work is a view of the zombie's mental life as being like a giant iceberg submerged under water, with only a fraction of the whole discernible above the surface in the shape of phenomenally conscious imagery, such that the zombie must infer the nature of the rest that's still under the surface. But this view is simply confused. If our sense that the zombie's experience is deeply impoverished relies on this view of self-knowledge, then we should reject the Liberal conclusion Horgan draws from these cases.

However, the Liberal can reply that access consciousness is one thing, and phenomenal consciousness something else altogether, and she is not here guilty of confusing the two. Rather, the Liberal claim is that the partial-zombie who lacks my agentive-cognitive experience would be missing a feature of my *phenomenal* consciousness. We can grant that the zombie would not find himself dependent on drawing inferences from his sensory experiences to decipher his own motivations for his behavior. But we still might want to claim that the zombie's self-knowledge is somehow deficient. The problem is *not* that the zombie couldn't be aware of what is going on inside, but rather simply that so much of his mental life *would be non-phenomenal*.

The point is not that we normal subjects have more self-knowledge than the zombie has, but that we have more of the right sort of self-knowledge. We have more phenomenal experiences, which afford us an immediate, acquaintance-like awareness of the contents of our thoughts, which in turn makes manifest in consciousness the reasons why we are acting the way we are and why we are enjoying the sensory imagery we are. The zombie might *know* all of these things as well—and perhaps know them non-inferentially—but he does not know them the way we do: in the way that makes them *flow through our awareness as part of our stream of consciousness*. This line of response is still open for the Liberal to take.

But of course, the Conservative can respond by denying that we *do* have agentive-cognitive phenomenal experiences that disclose the contents of our cognitive states in this special way, and thus deny that our self-knowledge will differ in kind at all from the partial zombie's. I do not see how this particular debate can progress much further beyond this point. Either we have the sense that these zombies are imaginable and phenomenally

distinct from us, or we do not—or we remain unsure whether we can tell either way based on introspection. So, it seems the appeal to agentive-cognitive phenomenology has not moved the dialectic between Liberals and Conservatives on much from where we were after examining Siewert's phenomenal contrast cases.

3. Which versions of the CP Thesis does Horgan's zombie contrast case argument support? Like Siewert, Horgan does not distinguish between Qualitative and Modal Theses, and seems to take his argument as offering support for both of them. On the one hand, he stipulates that the partial zombies that lack agentive-cognitive phenomenology are sensory phenomenal duplicates of us, so if there really is a phenomenal contrast between their experience and ours, the difference must be due to *non-sensory* phenomenology (the Qualitative Thesis). On the other hand, agentive-cognitive experience is supposed to be characterized by experiences of acting in accord with one's occurrent beliefs, desires, and purposes, such that the phenomenology is at least partly constituted by the conscious awareness of these thoughts—which of course necessitates that one is having thoughts (the Modal Thesis).

Suppose we do seem to discern a phenomenal difference between e.g. my partial-zombie's experience—described entirely in terms of sensory-perceptual experience and imagery—of hearing and responding to my wife's request for a carton of milk, and my actual experience in that same situation. Then, to the degree that we can be confident the description of the zombie's experience exhaustively lists all the sensory phenomenology a subject might have in that situation, the contrast case gives us some reason to accept the Qualitative Thesis (provided the phenomenology the zombie lacks is suitably associated

with thought). But once again, the mere fact that we are aware of a phenomenal contrast cannot establish anything on its own about the modal properties of the phenomenology in question. So Horgan must build additional claims into his argument, if he wants it to lead us to the Modal Thesis. He does so by claiming that without agentive-cognitive phenomenology the zombie would utterly lack awareness of his ultimate motives for his own bodily movements—they would seem “spontaneous” and “out of the blue” to him. In other words, as we have seen, Horgan implicitly appeals to a specific view of the relationship between phenomenology and self-knowledge, without which his contrast cases would fail to have implications for the Modal Thesis. Since this view of self-knowledge is dubious, Horgan’s argument provides only dubious support for Strong CP.

Nevertheless, I think consideration of Horgan’s zombie cases is instructive. Our discussion of them has illustrated that the appeal of at least some of the phenomenal contrast cases that philosophers use to argue for CP is dependent upon a certain view of what is in our streams of consciousness and the relationship between phenomenology and our awareness of our own thoughts. Indeed, as I will argue in the final section, this reliance on a certain (often implicit) view of our first-person access to our thoughts connects the Phenomenal Contrast arguments to the other major type of Liberal argument in the literature, the Self-Knowledge Argument. I suspect that Liberals find these contrast cases persuasive only because they also find the considerations that support the Self-Knowledge Argument persuasive as well.

V. Phenomenal Contrasts and Self Knowledge

1. The canonical version of the Self-Knowledge argument for CP comes from David Pitt (2004, 2011). Since a number of philosophers have ably responded to this argument, pinpointing its weaknesses, my exposition of the argument and objections to it will be derivative and brief.⁷⁵ I will then turn to consider the relationship between this argument and phenomenal contrasts.

Paraphrasing Declan Smithies (2013b, 11), we can reconstruct and simplify Pitt's argument as follows:

- P1** We have introspective (i.e. privileged, non-inferential, immediate, first-person) knowledge of the intentional contents of our thoughts
- P2** We have introspective knowledge of the intentional contents of our thoughts only if at least some of our thoughts have cognitive phenomenal character in the sense of the Modal Thesis
- C** At least some of our thoughts have CP in the sense of the Modal Thesis

Since denying the claim that we have introspective access to our own thoughts is intuitively very unattractive (though see Carruthers 2011), the Conservative's best hope of refuting Pitt's argument hinges on her response to its second premise. The intuitive idea behind P2 is that CP allows us introspective knowledge of our thought contents because direct awareness of these contents is constitutive of, or identical to, awareness of what it's like to undergo CP experiences. As Smithies puts it, "introspective knowledge of the phenomenal properties of experience provides us with introspective knowledge of its intentional properties" (2013a). This idea is in line with a general approach to explaining

⁷⁵ For discussion, see Byrne 2005, 2008; Levine 2011; Pitt 2004, 2011; Smithies 2013b; and Tye & Wright 2011.

introspective knowledge of our mental states: we have direct, non-inferential knowledge of what conscious states we are presently in only in virtue of our privileged access to their various distinctive phenomenal characters, i.e. what it is like for us to be in them.

As several philosophers have pointed out, we can straightforwardly dismiss P2 by appealing to any account of introspective knowledge that does not rely on phenomenal character to explain how we have privileged access to our mental states. Probably the most popular alternative is to posit some sort of reliable cognitive mechanism for forming introspective beliefs about what is currently going on in our minds. For instance, on one suggestion, “what it is to have knowledge of what one is thinking is to token a mental representation—a mentalese sentence—that expresses the fact that one is thinking what one is thinking. What makes this [introspective knowledge] is the fact that this sentence tokening is not the result of an inferential process, but rather an immediate causal result of the first-order thought state itself (together with some functionally characterizable internal monitoring process)” (Levine 2011). So CP is not necessary in order for us to have introspective knowledge of what we are thinking.

Pitt objects to this type of response to his argument, because he contends that our introspective knowledge of our thought contents involves an acquaintance-like awareness of these contents, just like our awareness of the objects of our perceptual states, or of what it’s like to undergo a sensory experience (2004, 19). But, as we saw earlier in our discussion of Horgan’s zombie cases, many philosophers will be only too happy to deny that we have *this* kind of awareness of our thoughts. So the problem for the Self Knowledge Argument is that it places the Liberal in a weak dialectical position, since the Conservative

easily can reject the background view of introspection that the argument needs in order to get off the ground.⁷⁶

2. However, while I think the Self Knowledge Argument is ultimately unsuccessful in establishing the Modal Thesis and the existence of Strong CP, the argument does seem to have one clear advantage over the various Phenomenal Contrast arguments, in that it marks out a fairly clear route from the truth of its premises to the truth of the Modal Thesis. For if the only way for us to have the kind of introspective access to our thoughts we manifestly do possess were for CP experiences to be constituted by acquaintance-like awareness of the contents of our thoughts, then introspective evidence would directly support the conclusion that the instantiation of certain phenomenal properties necessitates the entertaining of thoughts (on the assumption that such awareness of thought contents is or implies thought).

Now, I want to end the chapter by drawing a connection between the Self Knowledge Argument and the Phenomenal Contrast Argument, by showing how the former—or at least the intuitive view underlying it—is used, often implicitly, to supplement the latter. To do so I will try to demonstrate how Liberal proponents of the contrast case arguments often seem to be relying on something like the view of introspective knowledge I just outlined in their interpretations of the phenomenal contrast cases. In particular, I think this point helps explain how the proponents of these arguments

⁷⁶ However, see Smithies (2013b) for a (partial) defense of Pitt's view. In defending Pitt's claim that phenomenal consciousness is necessary for self-knowledge, Smithies relies on a similarly tendentious claim about the inadequacy of reliabilist accounts of knowledge in general, and self-knowledge in particular. So I don't think his defense succeeds in placing the argument in a much stronger dialectical position, but unfortunately I cannot address his points in detail here.

are able to reach the conclusions they do regarding the Modal Thesis, even though all their arguments begin only with a simple observation of a phenomenal difference between two hypothetical experiences.

I have already described how I think this works in the case of Horgan's zombie argument. I tried to make vivid the degree to which the motivation for Horgan's conclusion—that agential-cognitive phenomenology is ubiquitous in our mental lives—is reliant on a certain view he seems to assume about the nature of our access to our own occurrent cognitive states. The view essentially resembles Pitt's picture of self-knowledge as being dependent on phenomenal awareness. Horgan tries to persuade us that if our experience were entirely exhausted by sensory phenomenology, our awareness of our own inner lives would be radically impoverished. We would be just like a zombie whose own beliefs, desires, and reasons for acting are entirely opaque to it. And his argument needs this move in order for it to have any chance of supporting the Modal Thesis. For merely pointing to a phenomenal contrast between, on the one hand, a subject who experiences her bodily actions entirely in terms of sensory phenomenology, and, on the other, ourselves, whose agential experience of our bodily actions is apparently much richer, by itself cannot tell us about any metaphysically necessary connections between our richer agential phenomenology and thought. The initial datum is simply that when we imagine being the sensory-only zombie, we realize what it is like for it is somehow impoverished in comparison to our normal experience. This only gets us as far as the Qualitative Thesis.

The interesting thing is that Siewert makes a similar move in arguing for his conclusion that our understanding phenomenology is "inseparable" from occurrent awareness of the contents of what we are understanding. He considers an alternative

account of the contrast cases in which the contrast is explained by “a *non-sensory* experience of understanding—a perhaps more or less homogeneous ‘understanding-feel’—separable from the varying ways of understanding one’s words as one speaks” (2011, 21).

Siewert’s objection to this view starts from the claim that “we wouldn’t know we were reading with understanding in a given case, if what it was like for us to read then were no different than what it is like for us to read without comprehension: we’d think we didn’t understand” (22). In other words, if our conscious experience consisted entirely of sensory phenomenology (which we could enjoy without understanding), then we would have no way of knowing what, if anything, we were understanding a passage to mean, when we read with comprehension. But if CP just consisted of a uniform “understanding feel,” that feeling “at most could yield the mere knowledge you are *understanding something somehow*,” not awareness of the specific contents we were entertaining. In either case, we would be like one of Horgan’s partial zombies, entirely ignorant of much of the contents of our inner states. So in order to adequately explain our introspective knowledge of what we are understanding when we read with understanding in a phenomenal contrast case, we need to posit a Strong CP that necessitates awareness of thought contents and thus is modally bound to the actual exercise of thought. Thus, Siewert arrives at the Modal Thesis, having started his argument merely by observing a phenomenal contrast. But this line of reasoning clearly assumes something like Pitt’s view of introspective knowledge, according to which we know the contents of our cognitive states only by directly knowing what it’s like, phenomenally, to be in them.⁷⁷

⁷⁷ Siewert’s argument does not depend on the claim that other forms of self-knowledge are impossible or necessarily inadequate. He seems to allow that, perhaps, we *could* know the contents of our thoughts even if reading with and without comprehension were phenomenally the same. He only claims that humans *in fact*

What about our Siegel-inspired argument based on Perceptual-Recognition phenomenal contrasts? Recall that earlier I pinpointed the third premise as the initial moment where the argument shifts from consideration of phenomenal differences in particular hypothetical cases, to making a claim about certain properties that the phenomenology in question must have necessarily:

P3 The phenomenal properties of perceptual Before and After experiences supervene (with metaphysical necessity) on their intentional properties: if there is a phenomenological difference between these two experiences, then they must differ in intentional content

From this premise, the argument proceeds (roughly) by contending that the intentional content in question must be conceptual and thought-like, and thus that the phenomenological difference must involve a difference in the exercise of thought. I suggested that in order for this argument to succeed in reaching a Modal Thesis conclusion, its proponent must explain how she arrives at a claim about a necessary connection between phenomenology and intentionality simply by consideration of a handful of examples of phenomenal contrast.

The relationship between Siegel's argument for P3 and a Pitt-like view of self-knowledge is perhaps more subtle than in the cases of Horgan and Siewert's arguments, but nevertheless I think a similar connection can be drawn. Siegel argues that we should accept P3 because the phenomenology that accounts for the difference between Before and After experiences, which she assumes is some kind of feeling of recognition or familiarity, must involve representing *what* seems familiar. For "one would expect a raw feeling of

have phenomenally-mediated self-knowledge, on the grounds that otherwise our phenomenology of thought and understanding would be "epistemically superfluous," which, he thinks, we know it isn't. So his argument does not depend on the Self-Knowledge Argument, as formulated at the beginning of this section, so much as on the phenomenal view of self-knowledge underlying that argument.

familiarity [i.e. one which does not represent anything in particular as familiar], if there was such a thing, to leave one with a sense of confusion, since if it was clear to the subject what was being felt to be familiar, then this would seem to make the feeling representation after all" (498). Since we are *not* typically confused in situations where we experience perceptual recognition, the familiarity phenomenology cannot be "raw," but rather brings with it a representational change.

But once again, the argument seems to operate under a certain assumption, namely that if the contents of our thoughts are not made manifest to us in the phenomenal character of conscious experience, we inevitably remain ignorant of them. In this case, we are to imagine that the person who experienced this "undirected" familiarity feeling inevitably would be at a complete loss as to the identity of the *object* of familiarity, just as Horgan's zombie was at a complete loss as to the reasons he was moving his limbs the way he was. But Siegel has said nothing to rule out the possibility of there being cognitive mechanisms that direct a subject's attention to a salient feature in her perceptual environment and produce *non-phenomenal* states that communicate that the feature has been encountered before (or something to that effect: perhaps just that the feature is worthy of attention, is the object of her felt reaction, etc). And such non-phenomenal states could provide the requisite additional self-knowledge alongside the raw feeling. So it appears the default model of introspection Siegel is working with, as evidenced by what she overlooks in her argument for P3, is also something like the view that we only have immediate access to the contents of our various mental states via phenomenal character.

Finally, to take an example from a philosopher we have not considered in this chapter, Galen Strawson recently has offered what he calls his "Argument from

Interestingness,” which begins with the premise that “if cognitive experience [i.e. CP] didn’t exist, life—experience—would be pretty boring” (2011, 299). Since life isn’t that boring, but is “intricately and complexly interesting and various,” we can be confident that CP forms part of our conscious lives. But the reason Strawson gives for supposing that the sensory-only life would be boring is that the contents of our thoughts would not be *experienced*, meaning they would not be made manifest to us in phenomenology. Strawson allows that on the sensory-only view, there could be a generic “interestingness feeling” that accompanies our thoughts, but he thinks on this view, “although three equally fascinating propositions p, q, and r, may be very different, it’s their content that differentiates them, and this isn’t something that is consciously experienced in any way” (300). In other words, a large part of what prevents life from being boring is the fascinating nature of the countless propositional contents we might entertain, but unless we can somehow be made aware of these contents via our introspective awareness of phenomenal properties, we will miss out on them altogether. We might process them sub-personally, but “so far as phenomenological goings-on are concerned,” such processing is “no different from the way in which machines can be said to register conceptual content” (300). Since our inner lives are infinitely more interesting than machines’, we must posit the existence of CP that necessarily brings with it awareness of thought contents.

However, Strawson’s argument will only resonate with those who, like him, assume that genuine *personal-level* access to our thought contents must be clothed in *phenomenal* garb. Those philosophers who reject Pitt’s Self-Knowledge Argument—perhaps by appealing to a reliabilist account of introspection without phenomenology—are likely to reject Strawson’s Interestingness argument, and for the same reason. Strawson complains

that his opponents “can’t say that when we read there’s an experiential [i.e. phenomenal] difference between Hamlet-content-flavoured interestedness and excitement and Othello-content-flavoured interestedness and excitement, or Catch 22-content-flavoured interestedness and excitement. There are no such experiential flavours on their view. At the phenomenological level these works are significant only as powerful causes of interestedness-feelings, sorrow-feelings, and excitement-feelings of greater or lesser degree (apart from being different as sounds or marks)” (301). All of which is true, about the *phenomenology* that his opponents acknowledge. But it does not follow from this point that there are not differences in our *personal reasons* for being interested, sorrowful, and excited in reading each work, reasons of which we are intimately, non-inferentially aware as we read. It does not follow, unless we accept something like Pitt’s view of privileged self-knowledge.

I should underline that I am *not* arguing here that the view of self-knowledge driving the Liberals’ intuitions is false. I am pointing out only that this view easily can be rejected by many contemporary philosophers, Conservatives and otherwise, and that the various contrast arguments rely on it in order to transition from the evidence for non-sensory phenomenal differences to a Modal Thesis conclusion. These points explain why the Liberals’ case for Strong CP based on phenomenal contrasts is ultimately unpersuasive.

VI. Conclusion

I have endeavored to do two things in this chapter. First, I have tried to show that the Liberals’ case for CP, based on the Phenomenal Contrast and Self-Knowledge Arguments, at best only supports the Qualitative Thesis, but not the Modal Thesis.

Different types of Phenomenal Contrast Argument succeed to varying degrees at providing evidence for irreducibly non-sensory phenomenal properties, which plausibly are associated in some significant way with the faculty of thought. However, none of the phenomenal contrasts that we considered, on their own, provide evidence for phenomenal properties that vindicate the Modal Thesis. The Self-Knowledge Argument could support the Modal Thesis if it were successful, but it is undermined by its reliance on a certain controversial view of introspection that the Conservative is likely to reject. Thus, the Liberal needs to do a better job of defending this view before the Self-Knowledge Argument can be deemed a success.

The second thing I tried to accomplish in this chapter was to tease out an important connection between the two main Liberal arguments: certain Liberal formulations of the Phenomenal Contrast Argument implicitly rely on the very same view of introspective self-knowledge at the heart of the Self-Knowledge Argument, in their attempts to establish the Modal Thesis. Recognizing this fact allows us to understand one reason why the Liberal's interpretation of the contrast cases, and what they take these cases to show, so often differs markedly from how the Conservative understands the significance of these cases. Further, it allows us to see why the contrast arguments for the Modal Thesis are ultimately unsuccessful.

Chapter 4

Cognitive Phenomenology and the Fringe

I. Introduction

1. The burden of the last chapter was to show that nothing in the recent literature on cognitive phenomenology succeeds in providing significant support for the existence of what I have called Strong CP, the kind of cognitive phenomenology whose existence entails both the Qualitative Thesis and the Modal Thesis. Since my primary target in this thesis is the Modal Thesis, the claim that there is a kind of phenomenal property whose instantiation metaphysically necessitates the thinking of a thought, it was important, first,

to establish that a compelling case for the Modal Thesis has not yet been offered, and, in that sense, to side with the Conservatives against the proponents of Strong CP.⁷⁸

However, there is a *hard-line* and a *concessive* approach one could take towards the defenders of Strong CP. The hard-line approach is to deny that the Liberal arguments have any real merit, or at least to insist that we should remain Conservatives by default, as the burden of proof lies with those trying to convince us otherwise and they have failed to meet it.

The more concessive response, which I will continue to develop in this chapter, elaborates and defends the existence of a kind of Moderate CP, which offers an alternative explanation of the sorts of considerations that Liberals point to as evidence for Strong CP. I will argue that there is a group of phenomenal properties that are specially associated with cognition and qualitatively distinct from sensory phenomenal properties, but that the instantiation of these properties in a subject does not necessitate contentful thought. I will show how acknowledging the presence of this phenomenology can explain the intuitive appeal of the popular Phenomenal Contrast arguments at least as well as positing Strong CP can, and enjoys considerable plausibility in its own right. Thus, chapters 3 and 4 together form a one-two punch against those who appeal to the standard Liberal arguments to support a view of CP that entails the Modal Thesis.

The Moderate CP account I develop should be of interest to open-minded people on both sides of the cognitive phenomenology debate and to all those on the fence, insofar as I

⁷⁸ Of course, one could defend the Modal Thesis independently of Strong CP, by rejecting the Qualitative Thesis. But as I pointed out in chapter 1, this view is a version of empiricism that, as far as I know, is not accepted by any contemporary Phenomenal Liberals.

am attempting to offer a plausible and genuinely distinct middle way between the two extremes.

2. I first want to motivate my Moderate CP view with an example. In his defense of cognitive phenomenology (2012), Anders Nes asks us to compare the following two stories:

The story of Thoughtful: as Thoughtful sits at his desk one morning, a friend comes in, saying ‘Have you heard? A BA plane has just crash-landed at Heathrow.’ Oh gosh, Thoughtful thinks to himself, a BA plane has just crash-landed at Heathrow. As he does so, he subvocally repeats his friend’s words to himself (letting the thought sink in, as one might say). He is naturally surprised to hear this news. His latent fear of flying is reactivated in a surge of anxiety about his upcoming BA flight into Heathrow. He is tormented by unpleasant visual imagery as of a BA plane lying broken in pieces on the grass beside the Heathrow runway.

The story of Thoughtless: Thoughtless sits in a qualitatively identical office; his friend comes in and says just the same words. Although Thoughtless has enough of a sense of English to parse the sentence uttered, he is in no position to understand what it means....Nevertheless, he subvocally repeats his friend’s words to himself. For some reason, he has a reaction of surprise at something or other, and a latent fear of something or other flashes up as a surge of anxiety in his mind. He has visual imagery as of segments of an approximately cylindrically shaped body irregularly placed on a green flat surface with a greyish rectangular field. He finds, for some reason, this imagery unpleasant. The phenomenology of his sensory experiences, his subvocal sayings, his visual imagery, his feeling of surprise and his surge of anxiety are as similar as they can be to their counterparts in Thoughtful’s case, consistently with the absence in his case of any thought of planes, crashes, airports or Heathrow, none of which flows through his mind. (90)

The conclusion that Nes hopes we draw from the comparison is that there is a phenomenal feature of Thoughtful’s experience that is missing in Thoughtless’s case: all of Thoughtful’s sensory and emotional experiences “circle around” the thought about the plane crash. As Nes explains, “it is... an aspect of Thoughtful’s phenomenology that the properties visually imagined are the properties of the same sort of event—a plane crash—as the words are experienced as speaking of, and as is experienced as surprising and frightening.” (91).

The same cannot be said of Thoughtless’s experience. After all, “all we know is that

[his sensory and emotional experiences] occur at pretty much the same time, in the mind of the same person,” and, “clearly, it is possible for people to be anxious about one thing while they talk to themselves about another wholly unrelated thing and entertain imagery about yet a third thing wholly unrelated to either.” So, it seems that “it is not an aspect of Thoughtless’ phenomenology that what is visually imagined as being so is the same sort of event, or indeed has anything to do with, what is found surprising, or what is found anxiety-evoking, or what the words uttered are speaking of” (91-92).

Nes takes his stories to show that thoughts play a constitutive role in conscious experiences, as the “unifying themes” around which our other experiences are felt to revolve. In other words, what Thoughtless lacks is Strong CP—phenomenology that implicates thought content—to unify his imagistic and emotional experiences. But it seems to me that the phenomenon to which the stories most obviously point is not the role that the thought content *per se* plays in unifying the experiences, but just the fact that *our experiences of thinking seem to be thematically unified at all*. Nes has not given us good reason to think the phenomenal contrast between Thoughtful and Thoughtless cannot be explained by positing an experience of our imagery and feelings being unified around a common topic *without* that topic itself being phenomenally conscious. Indeed, if asked to describe what I experience about the topic *itself*, it seems fair to describe just those sensory experiences that are unified. In short, there is space for someone to grant the claim that Thoughtless lacks an experience of thematic unity that we enjoy, but deny that this

experience entails that thought contents are explicitly spelled out in phenomenal consciousness.⁷⁹

So the intuitive puzzle this case raises is *if all phenomenology is sensory-perceptual, how are my perceptual, imagistic, and emotional experiences ever experienced as relating to a single topic at all, rather than as entirely unrelated?* And the most straightforward solution to this puzzle is just to posit a distinctive type of non-sensory experience—which Thoughtful has and Thoughtless lacks—of these various sensory experiences being somehow *related* to the contents of the subject's present thoughts.

3. There are two things I wish to highlight about this solution. The first is the way it can be generalized to provide a response to the questions raised by other phenomenal contrast cases, including those we discussed in the last chapter. For instance, when considering the alleged difference between normal subjects and Horgan's sensory-only zombies, it is natural to suppose that a normal subject's actions are experienced as purposeful, because they are experienced as somehow *connected* to her beliefs, desires, and intentions. Similarly, Siewert's contrast case of reading with and without understanding reveals the sense we have when we read comprehendingly that the words on the page are determinately meaningful for us. A plausible explanation of this type of experience is that the words seen on the page (or heard spoken) are experienced as relating to—perhaps as

⁷⁹ One might argue that the contents, the "themes," must themselves be conscious since otherwise Thoughtless, with only sensory experience to draw on, would necessarily be ignorant of the meaning and significance of his own sensory experiences. Indeed, Nes encourages this view when he describes Thoughtless as not understanding his friend's words, and as having images and feelings about "something or other," about which he is, apparently, clueless. But, as I argued at length in the last chapter, the presence of the right (phenomenally unconscious) cognitive states should be sufficient to grant a subject like Thoughtless *understanding* of language he speaks or hears and introspective *knowledge* of his own thoughts. At least, to say otherwise is to make a contentious claim.

prompting, or as “standing in for”—the particular contents of the thoughts we are entertaining at that moment. An analogous account can be given for how we experience our inner speech and other sensory imagery, not just as vague “pictures” aimlessly floating through consciousness, but as having a precise, intelligible meaning beyond their surface qualitative properties: we experience the images as “proxies” for the determinate contents of our thoughts.⁸⁰

In each of these examples, it is hard to capture the phenomenal element that seems to be missing in an entirely sensory description of our conscious lives—a description like those of my partial zombie and Thoughtless—without making reference to some sense, some awareness, of the meaning that these sensory experiences have for the subject. Each case illustrates that sensory experience seems to be shot through with a sense of meaning-for-the-subject, a meaning that goes beyond the information immediately conveyed by sensory phenomenology alone. And one way to understand this semantic experience is as a conscious awareness of our own ongoing, personal-level cognitive activity. So these cases motivate the idea that we need to posit *some* kind of awareness or representation of thought within phenomenal consciousness, quite apart from the issue of whether we acknowledge specific types of phenomenology that somehow implicate the *contents* of

⁸⁰ In a well known passage in the *Philosophical Investigations*, Wittgenstein considers a picture of an old man walking up a hill with a walking stick, and points out that, so far as the picture’s details are concerned, the picture could just as easily be interpreted as depicting the man sliding *down* the hill, digging his stick into the ground to slow his descent (1968, 139-140). Wittgenstein’s point is about the inadequacy of the intrinsic properties of images, mental or otherwise, to fix reference, or as he says, to “force a particular application on us” (140). But this raises the question of how it can *seem* to us that our mental imagery has a determinate and transparent meaning. If we think, for instance, of London Bridge, we are likely to experience a fleeting visual image of a bridge, which might just as well be used to represent any number of bridges (it might resemble the Golden Gate Bridge more than London Bridge, for example). Yet, arguably, we experience the image as depicting what we are thinking about, namely this one particular bridge. We experience the image as having the unambiguous meaning we intend it to have.

thought. In other words, we at least should posit a form of Moderate CP to explain these types of cases.

The second thing I want to point out about this solution is how it occupies a middle ground between the common Liberal and Conservative responses to the puzzle posed by phenomenal contrast cases, which we considered in detail in Chapter 3. The Liberal response is to seek to explain the extra phenomenal feature that seems to be missing from sensory-only accounts in terms of Strong CP: we experience our actions as purposeful because we experience them as following on from specific intentions whose content is phenomenally conscious; we experience our sensory and emotional experiences as “revolving around” a thought in virtue of experiencing the content of that thought; and so on.

But the challenge for the Liberal is in clearly distinguishing between what conclusions about CP the examples directly support from introspection alone, and what conclusions they only support, at best, by an inference to the best explanation. We might experience our actions as following on from intentions without directly experiencing those intentions, experience our imagery as determinately meaningful without directly experiencing its determinate meanings, and experience our emotional and sensory episodes as being about a common subject matter without directly experiencing that subject matter.

If the Liberal’s response is that introspecting on these experiences is enough to rule out this possibility, fair enough. But then he is placing a lot of weight on the deliverances of introspection, and the Conservative is simply going to deny the claim. We saw in Chapter 3 that phenomenal contrasts put some degree of pressure on us to accept the Qualitative

Thesis: at least some experiences of thinking seem to involve phenomenology beyond the sensory. That there is *some* sort of, perhaps entirely *general*, non-sensory phenomenology of meaning that characterizes our experience in each of the above examples is a less ambitious claim, and is more difficult to deny, than the claim that the contrast cases manifest differences in Strong CP. It depends only on taking the contrast cases as evidence for genuine phenomenal differences in these cases, and on an inclination to characterize what these phenomenal differences are like in terms of the cognitive-semantic properties of the subjects' sensory experiences.

On the other hand, the Conservative response to these cases is to explain away the apparent phenomena—the sense of purpose, meaningfulness, and unity that seems to characterize our sensory-perceptual experiences—entirely in terms of the influence of non-conscious cognitive states. On this view, we do not, strictly speaking, *experience* the cognitive-semantic properties associated with sensory features of our stream of consciousness, if 'experience' is taken to imply enjoying phenomenal character. Rather, it is natural for us to judge that our sensory experiences “seem” to exhibit these characteristics, provided we have the right sort of access to those cognitive states that function to make our actions purposeful, our imagery determinate and intelligible, and our conscious stream coherently organized—and this is possible even if the cognitive states are not clothed in phenomenology at all. We err when we think that the imagined subject who only enjoys sensory experiences would lack some features of our conscious experience, since what we imagine as the missing phenomenal features is really just the (non-phenomenal) impression of “things running smoothly upstairs” that one gets when one's experiences and

unconscious cognitive states are properly functionally organized, and both ourselves and sensory-only “zombies” have this organization.

Of course, the Conservative is free to draw this conclusion. But in doing so, she has to flatly deny that there are ever any genuine phenomenal contrasts described in the above scenarios. Note that this is a different strategy to the one the Conservative typically takes in responding to contrast cases. The Conservative usually acknowledges that there is *some* phenomenal contrast, which indicates *some* addition of phenomenal character accompanying a subject first coming to understand/entertain/realize some thought content. She just seeks to account for this phenomenal addition in terms of a change in sensory experience. But in the above examples, the alleged non-sensory phenomenal feature I described is unlikely to be mistaken for a sensory feature, since each represents a semantic property that modifies or characterizes sensory experiences in the stream of consciousness, and it is hard to see *other* sensory experiences playing this role. It is hard, for instance, to see how the addition of another sensory experience could make a cluster of sensory experiences seem unified around a common theme. So it seems safer for the Conservative to simply say that these semantic properties that characterize our sensory experience, though we can become aware of them, are not phenomenally conscious at all. But I think this should be seen as a fairly extreme position to take, particularly in light of some of the considerations raised in chapter 3, and those I will raise in the rest of this chapter.

In any event, there is a promising position to be marked out between the Liberal and Conservative responses, a position that grants there is a genuine experience described by these cases that is non-sensory and represents the relations between sensory experience

and occurrent thought, but which does not determine or disclose the contents of thought in consciousness.

It seems to me that in order to explain the sorts of phenomena that we want CP to explain, there must be two components to the representational structure of a Moderate CP experience. First, there must be that component of the overall experience that conveys the semantic features apparent in consciousness, the sense of genuinely meaningful cognitive activity in the subject's mind. The second component that must be present in the experience is the representation of the relationship *between* these semantic features, this meaningfulness, and the various sensory items in the stream of thought. In this way CP can serve as the "connective tissue" of consciousness and cognition. I think the account I am about to offer captures both these components of cognitive experience.

In the next section, I sketch a neo-Jamesian account of Moderate CP. In section III, I address a Conservative challenge to this account, viz. whether the phenomenology in question should really be considered to be a form of CP. In section IV, I address a Liberal challenge to this account, viz. whether such phenomenology can really explain all the phenomenal contrasts that CP is typically called upon to explain.

II. A Neo-Jamesian Account of Moderate CP

1. The Jamesian Fringe

The notion of cognitive phenomenology I have just described receives significant support from a cluster of ideas that have been fruitfully explored in the cognitive science literature for the last twenty years. In fact, these authors did not originate these ideas, but drew inspiration from the seminal work of William James in his *Principles of Psychology*

(1890). The principal purpose of the rest of the chapter is to bring these neo-Jamesian insights to bear on the philosophical debate about cognitive phenomenology.

The key passages in the *Principles* for our purposes are found in the well-known chapter on the “Stream of Thought.” Indeed, the famous central metaphor of the chapter, from which it gets its title, is especially relevant. James writes,

The definite images of traditional psychology form but *the very smallest part* of our minds as they actually live. The traditional psychology talks like one who should say that a river consists of nothing but pailsful, spoonsful, quartpotsful, and other molded forms of water. Even were the pails and the pots all actually standing in the stream, still between them the free water would continue to flow. It is just this free water of consciousness that psychologists resolutely overlook. Every definite image in the mind is steeped and dyed in the free water that flows around it. With it goes the sense of its relations, near and remote, the dying echo of whence it came to us, the drawing sense of whither it will lead. The significance, the value, of the image is all in this halo or penumbra that surrounds and escorts it, -- or rather that is fused into one with it and has become bone of its bone and flesh of its flesh. (255)

Beginning with Bruce Mangan (1991, 2001), a number of cognitive scientists have pointed out that James’s argument in this chapter calls attention to a fundamental but often missed fact about our everyday conscious experience.⁸¹ Of course, the image suggests that consciousness is stream-like: it seems to flow. But conscious experience “flows” only because of what James here calls the “free water of consciousness,” which “flows around,” “steeps,” and “dyes” the more obvious elements of consciousness. The point of the idea that consciousness is like a stream, then, is not primarily that it is smooth or fluid *per se*, but that it contains two fundamentally distinct types of experience. The more apparent “pails and pots” in the stream include perceptual experiences and sensory imagery, what James calls the “nucleus” of thought. These nuclei are to be contrasted with the more subtle and evanescent waters of the stream, which for James encompass a great many features of

⁸¹ For instance, see Epstein 2000, 2004, and Galin 1994.

consciousness—making up the vast majority of the conscious mind, if we take what he says in the above passage seriously—but which he particularly identifies with a “sense of relations” that gives “significance” and “value” to the experiences in the nucleus. He claims that these kinds of experience make up a “fringe” of consciousness, distinct from, but ever surrounding, the nuclear sensory images.⁸² Furthermore, this distinction involves a difference in the qualitative properties of each type of experience: the fringe, in contrast to the nucleus, is phenomenally *non-sensory*.⁸³

We should avoid thinking that the distinction between nuclear and fringe experiences is essentially the same as the one commonly made between *attended* and *unattended* features of experience. Being in the fringe is *not* just a matter of being out of the “spotlight” of attention (a fact that James’s terms obscure). To see this, note that a fringe feeling that “permeates” a mental image is not necessarily unattended: we can be especially struck by the familiarity of a face we are looking at, or by the relevance of some detail of the visual scene to a thought we were having a moment earlier (say because the detail illustrates the thought especially well). In such cases, we attend to the face or the visual detail, but also to its familiarity or relevance. So, arguably, being unattended is not

⁸² For ease of style I will switch freely between speaking of a single nucleus and consecutive nuclei. However, I will always assume that there is only one nucleus in consciousness *at a given time*. So a single nucleus could be composed of a number of distinct sensory experiences, depending on how we individuate experiences.

⁸³ Note that, as suggested by the stream metaphor, the “waters” of fringe experiences both *extend between* consecutive nuclei as well as *surround and permeate* the current nucleus. In other words, they relate to the nucleus both diachronically and synchronically. This helps to clarify what sorts of experience James had in mind in speaking of the fringe. For instance, the senses of “relations” and “tendency,” which both reach backward in memory and forward with a feeling of anticipation, belong to the fringe feelings that occur between successive sensory experiences in the nucleus. It is from these feelings that consciousness can be said to receive its stream-like character, as they make for smooth transitions between thoughts. On the other hand, the fringe should be seen to include phenomenal properties that co-occur with and serve to modify sensory experiences in the present nucleus, such as when one is struck by the sense that a certain face in a crowd is salient or familiar. These are the features of the fringe that are more naturally seen as indicating the “significance” and “value” of an image or percept in the nucleus.

necessary for an experience to count as part of the fringe. Further, being unattended is not sufficient for an experience to be part of the fringe, as is clear from considering the constant stream of sensory experiences that lie at the periphery of consciousness and go largely or wholly unattended.

Still, I think it's safe to say that sensory items in the stream of consciousness are typically more easily attended to, whereas non-sensory fringe feelings often lie outside of our attention or direct introspective awareness. This difference naturally follows from the difference in the *types of information* that fringe and nucleus experiences carry: the sensory experiences in the nucleus make us aware of important features of our external environment or else highlight distinguishing features of the content of our occurrent thoughts (as a kind of "shorthand" for this content); by contrast, fringe experiences summarize contextual information *about* the items in the nucleus, e.g. about relations between the information in the nucleus and additional, presently unconscious information. So we should not be surprised if attention is more naturally drawn to nuclear information, as the more appropriate subject matter of our conscious interest.

Since fringe experiences are supposed to be non-sensory, their phenomenology is eligible to count as a type of Moderate CP, though we will return to this issue below. Furthermore, my earlier description of Moderate CP as the "connective tissue" of consciousness seems to fit perfectly with James's idea of the fringe as tying a current thought to earlier and later thoughts by signaling a "dying echo" of what came before and a "drawing sense" of what is to come.

At this point, some readers might be doubtful as to whether there exists such a thing as fringe consciousness. In response, we can point to several especially vivid examples of

it, instances in which the fringe seems to step out from the shadows of the sensory and into the foreground of consciousness, so to speak. My hope is that the reader will not have too much trouble in realizing how the fringe extends beyond these cases, to the sorts of more subtle experiences that can explain the phenomenal contrast cases.

2. Cognitive Feelings

Several types of fringe experience have already been acknowledged and studied separately within psychology. Psychologists have related these experiences, so-called “cognitive feelings,” to a large and diverse body of research, of which the fringe literature is only a modest part.⁸⁴ Russell Epstein, summarizing James, has compiled a helpful (non-exhaustive) list of prominent experiences found in the fringe, what I will call “fringe experiences,” “fringe feelings,” or “cognitive feelings”:

They include (1) the feeling of expectation we have when our attention has been drawn to something and we have some sense of what it might be, but we have not yet determined exactly what it is; (2) the feeling of having a word “on the tip of one’s tongue”; (3) the sense we have when we know that something is familiar; (4) the sense of connection provided by words that connote the logical structure of a thought, such as “but” or “nevertheless”; (5) the feeling of intending to say something, when we have a perspective of a scheme of thought that is not yet articulate; (6) the sense we have for the overall scheme or form of a book, a work of art, or a scientific system; (7) the sense we have that there is more to a thought, even when we do not have the words to complete it; (8) the sense of anticipation that occasionally causes us to mistakenly use an anticipated word instead of the correct one; and (9) the sense of being “on the right track” to a conclusion. (Epstein 2000, 553-554)

Let’s look at three of these examples in more detail.

⁸⁴ For instance, cognitive feelings have been examined in articles on dual process models of reasoning (e.g. Price & Norman 2008), intuitions (e.g. Topolinski & Strack 2009), metacognitive judgments (e.g. Koriath 2007), perceptual and conceptual fluency (e.g. Reber et al. 2004), and aesthetic and emotional experience (e.g. McGovern 1993; Epstein 2004), to name but a few.

First, imagine a situation in which you pass someone in the street and have a strong sense of recognition, which somehow seems directed upon this person's face, even though you cannot at present call to mind how you recognize them. Let us call this a feeling of familiarity (FoF). These experiences can be very vivid, such that I doubt many would deny that there is something it is like to have them.

A second kind of cognitive feeling has much in common with the FoF, but I think it is subtly different. Imagine sitting deep in thought, working through a philosophical question. You are mentally laying out, step by step, how you would present your position, and anticipating possible objections or promising implications. As your attention is focused on one point, and you are still occupied in fully developing its details in consciousness, you have a sense of something else: an objection or implication that follows from this point, but is as yet entirely inchoate, merely implied, "felt" but not "seen". You are anticipating something that follows from your current thought, but have no clear awareness of what that something is. In other words, what is in the fringe of consciousness at that exact moment is merely the information that there is this further thought, "waiting in the wings" until your attention has room to usher it in. It might be accompanied by a subtle positive or negative feeling, signaling that the non-conscious information is either agreeable or disagreeable. And at times there can be some fuzzy sensory experience, at the edge of attention, which serves to stand in for this still absent thought; but there need not be any such sensory accompaniment, and even when there is, it is only experienced as suggestive of the non-conscious information *because* it is experienced in conjunction with the cognitive feeling.

Call this cognitive feeling the feeling of “Imminence” (Mangan 2001). Several of the examples on Epstein’s list (1, 5, 7, 8) have this feeling, or something very much like it, as a component. It is also what James alludes to, in the stream of thought passage, when he speaks of “the drawing sense of whither [the current thought] will lead” (or, we might prefer to say, where it *might* lead, since at least some of the time we can choose whether to ignore or respond to the feeling).

Next, consider the cognitive feeling that accompanies a moment of sudden realization. The moment in question can involve remembering something, seeing the solution to a logic or math problem, or getting a joke. For instance, recall Siewert’s example of a friend walking up to him and asking if he “has the book” (2011). After initially having no idea what she is talking about, he suddenly remembers she had asked him to return a book he had borrowed. In such a moment, we might image the book’s title or its cover, but we would also experience an instant rush of certainty, where before there had been a sense of puzzlement. It is a feeling of “the puzzle pieces all falling into place,” what people sometimes describe as an “ah ha moment,” and what Mangan calls simply “Rightness” (2001). Rightness (along with its opposite, “Wrongness”) is described by James as “the most important feeling in these fringes...the mere feeling of harmony or discord, of a right or a wrong direction in the thought” (261). As he points out, this feeling of assurance can make even gibberish, or the bizarre scenarios of a dream, seem to make complete sense.⁸⁵

⁸⁵ Cf. Robinson 2011. This is not to say that the feeling of rightness always communicates certainty. I take example (9) on Epstein’s list above, the sense of being “on the right track,” to be a less intense version of the same experience. Indeed, there are many situations, for instance when doing philosophy, in which a very *faint* feeling of rightness leads our thoughts in a certain direction in the face of considerable uncertainty.

The three examples above are meant to illustrate the sorts of experiences I have in mind, but there are likely many other fringe experiences.⁸⁶ Indeed, while the examples on which I have focused are quite striking and unusual experiences—and thus, more difficult for the skeptic to ignore—the picture of the fringe I am advocating, following James and others, is of a family of experiences that also includes many types of experiences that are more integrated and commonplace, phenomenology that is barely noticeable but that colors much of our everyday thought life.

Still, while fringe experiences are likely to be many and diverse, I think we can discern commonalities between them. At least a great many fringe experiences—including our three examples and the others on Epstein's list—seem to exhibit one or both of two broad categories of fringe phenomenology. These two important phenomenal types correspond nicely to two things James says about the fringe in his initial description when he says that the fringe brings with it “the sense of [the nucleus's] relations, near and remote, the dying echo of whence it came to us, the drawing sense of whither it will lead. The significance, the value, of the image is all in this halo or penumbra that surrounds and escorts it” (255).

The first phenomenal type corresponds to what he calls the “senses of relation,” and serves to *locate* the present thought—or more precisely, the present sensory experiences that form the nucleus of consciousness and serve as placeholders for the thought—within a network of associated cognitive representations that have or could become conscious through imagery. We can call this “locating” fringe (L-fringe) phenomenology. The second

⁸⁶ For instance, we have not discussed the famous “tip-of-the-tongue” (ToT) experience, but it seems best to treat the ToT as a composite of several other fringe feelings, such as a feeling of knowing, followed by wrongness (the opposite of rightness) if wrong words are proposed, then imminence, and then finally, if the desired word is actually identified, the feeling of rightness.

broad type corresponds to what James calls the “significance” and “value” of the image, and communicates *evaluation* of what is being entertained in the nucleus, particularly with respect to the subject’s unconscious cognitive information. Call this “evaluative” fringe (E-fringe) phenomenology. So, for instance, the cognitive feeling we earlier called “Imminence” arguably can be broken down into a locating element—signaling the presence of some novel idea that is ready to be brought into consciousness via imagery—and an evaluative element—signaling that this new thought is somehow positive or negative, consistent or inconsistent with the present nucleus, clarifying or confusing, and so on.

I have tried to describe the phenomenology of fringe experiences in terms of relatively simple cognitive *feelings*, roughly on a par, in terms of complexity, with sensations or simple perceptual qualia. While these experiences are often (perhaps always) about, or directed at, other things found within consciousness, and in that way have an “aspectual” element (cf. Robinson 2011, 206), at least on the face of it their content need not be conceptual: you don’t need to have the concept of familiarity in order to enjoy an FoF, for instance.⁸⁷ And undergoing these experiences could come apart from actually thinking any thoughts: you could feel like you have something interesting to articulate, but it could turn out that you don’t.⁸⁸ So these cognitive feelings do not vindicate the Modal Thesis.

⁸⁷ This comment is not meant to suggest that these feelings are “raw feels” or “mental paint”—they may have content and may represent it necessarily. Now, whether the FoF, for example, literally represents the property of familiarity or not, is a hard question I don’t know how to answer.

⁸⁸ The idea that the fringe can provide the conscious subject with all sorts of information about her own thoughts without itself (a) *being* a form of thought, (b) *representing* particular (contents of) thoughts, or even (c) *employing concepts* for thought, thinking, etc, relates nicely to some comments by Price (2002). Comparing Mangan’s theory of the fringe with Koriat’s (2000, 2007) work on metacognitive feelings, he says, “Feeling of Rightness, or ‘fit’, can be thought of as a metacognitive representation of a particular state of affairs that holds between different knowledge structures. Metaphorically speaking, we are conscious of the relational linking arrows, but not the entities linked...FOK, TOT, familiarity, preference and so on all reflect

III. Why the Fringe Qualifies as Moderate CP

1. Since the Moderate CP view that appeals to fringe phenomenology is supposed to provide a middle road between the typical Liberal and Conservative positions, the view naturally faces objections from both directions. The most obvious Conservative response to this proposal is to dispute whether the fringe should really count as a form of CP at all. In order to defend the claim that the fringe qualifies as a version of CP according to the Qualitative Thesis, we need to show that fringe phenomenology really is phenomenally distinct from all sensory phenomenal properties, and that it is rightly associated with the faculty of thought.

First, in what way is the fringe associated with thought? In what sense is it rightly considered a genuinely *cognitive* phenomenology, a phenomenology *of* thought? Bruce Mangan makes the plausible suggestion that the functional role of the fringe in general is to “represent almost all context information in consciousness”—including stored representations associated with the ones being consciously entertained, and the subject’s present goals—and thereby to “help mediate the voluntary retrieval of new information into consciousness” (2001, 1-2). This insight helps to shed light on what all fringe experiences have in common, such that it makes sense to think of the fringe as a distinct, unified class of experiences that has a vital role to play within cognition. Just as perceptual experiences inform us about our external environment, the fringe, in tandem with sensory

knowledge about our mental processing. Awareness of shifts of visuo-spatial attention, proposed...above as an explanation of subjects’ ability to localise masked targets, can also be thought of as a form of metacognitive knowledge. In this example, there appears to be fringe awareness of an attention shift without access to the antecedents of the shift (i.e., the subliminally presented information).” (16-17).

imagery, informs us about what is happening inside our own minds, including about cognitive processing that occurs “below the surface” of conscious awareness but is related to our present thoughts.⁸⁹

For example, consider the three cognitive feelings we described earlier. Mangan suggests that the feeling of familiarity serves to “[signal] that something now in consciousness has been encountered before”—and this seems about right (2001, 3). In this way, FoF’s are very similar to another well-known cognitive feeling, the “feeling of knowing” (FoK), which represents that we know a particular fact, i.e. that a particular representation is stored in our memory, even if we cannot presently access the relevant information in consciousness.⁹⁰

Clearly, the feeling of Imminence has a central role to play in bridging the domains of conscious and unconscious cognition. This experience implies that there is much more information that is accessible to consciousness than is currently in the stream, and helps to guide the subject’s conscious retrieval of pieces of this information over time. Mangan makes the important observation that this kind of feeling helps the mind solve a serious design problem, viz. that of how to make large amounts of useful information available to the conscious subject in spite of the extremely limited processing capacity of consciousness. Consciousness is so taxing to our cognitive resources that we can only focus on a few bits of information, a few experiences, at a time. The challenge is to suggest to the

⁸⁹ I see no reason to think that this account of how we consciously experience cognitive context and the “reliable cognitive mechanism” model of our (non-phenomenal) introspective access to the content of our beliefs and occurrent thoughts, which we discussed at the end of the previous chapter, are mutually exclusive.

⁹⁰ It may be that the FoK and FoF are identical experiences, but I think it is more likely that one is a species of the other, or that they are simply related but distinct. One suggestion for how to distinguish them is that the FoF has to do with episodic memory, whereas the FoK has to do with semantic memory.

subject the presence of additional unconscious information without using up very much of the valuable resources of consciousness in doing so (2001, 14).⁹¹

The functional role of Rightness seems to be to tell us that things are going well in the cognitive domain, and things are going well in the cognitive domain when the information in the center of conscious awareness “fits” with the context created by the subject’s unconscious processing (cf. Mangan 2001, 4). The context involves not only information with which the conscious nucleus is associated, but also the current goals of the subject. For instance, when I remember the borrowed book about which my friend has asked me, the rather limited contents of the nucleus—imagery of the title and cover of the book, say—are suffused with a strong feeling of Rightness, because these contents fit both with the larger web of associated information connecting books, me, and my friend (e.g. memories of my friend giving me the book) *and* with a current behavioral goal, namely to make sense of the question my friend just asked me. As long as we know that there is this fit between nucleus and context, we know our thinking is on the right track, without needing the details of the context itself to enter consciousness.

If this picture is on the right track, the fringe is much more than the occasional queer feeling, an epiphenomenal oddity added to the truly important business of consciousness and cognition. Rather, in making us aware of information about our cognitive context, and how this information relates to our present thoughts, the fringe plays a crucial role in the monitoring and control of higher-level thought processes. We might say it helps direct the

⁹¹ Mangan offers the memorable analogy of a computer word processing program: in order to leave the majority of limited screen space for the words of your document and not clutter up the interface, the program hides most of its features in drop-down menus, which are represented on the screen as small icons that only open when clicked on. The feeling of Imminence and other fringe experiences are like these highly economical, inconspicuous icons.

course of the stream of thought. In other words, a straightforward way to see why the fringe counts as a phenomenology of thought is to recognize that it is a phenomenology that tells us *about* thought.

2. Still, even if she grants that the fringe exists and has this cognitive role to play, the Conservative might object that fringe phenomenology reduces to a type of sensory phenomenology. In response, I think we can offer two reasons for supposing that fringe phenomenology is qualitatively distinct from all sensory phenomenology.

First, we can appeal to the introspective evidence: the fringe *feels* cognitive, rather than sensory. Arguably, a prerequisite for all sensory experience is that it involves an experience of spatial location (even if the represented location is just *everywhere* or *all around the subject*), but, on the one hand, certain vivid cognitive feelings such as *déjà vu* do not seem to have any spatial location. And when fringe experiences do seem to be located in the head, they do not feel like bodily sensations (in spite of its name, the tip-of-the-tongue experience does not feel like a sensation on the tip of one's tongue).

On the other hand, at least a great many fringe feelings serve to modify, and thus are directed at, aspects of sensory experience that are spatially located. But it is hard to see how any distinct feature of sensory experience itself could perform this task. For instance, you could experience all the same sensory elements that make up your visual percept of the face before you, without thereby experiencing the face as familiar.⁹² And no strictly

⁹² In the last chapter, I criticized Siegel's claim that experiences of recognition or familiarity must represent an object of familiarity/recognition, a criticism that may seem to run counter to my position here. But my objection was not to the claim that our familiarity experiences always or typically involve an experienced object, but to her contention that if we did not experience such an object, as part of *phenomenal*

sensory addition—a shiver down the spine, a sensation of furrowing the brow—would seem to be adequate to characterize either the familiarity feeling itself or the sense of its being *about* the face.⁹³

In chapter 1 we said a necessary condition for a phenomenal property being sensory (i.e. a perceptual experience with low-level content) is that it be phenomenally basic, and the fact that many fringe phenomenal properties are experienced as modifying sensory features of experience—we experience the FoF as a quality of the visual features of the face, rather than just as an undirected feeling—provides at least a *prima facie* reason for thinking that these feelings cannot occur without the sensory experiences they modify, and thus are non-basic. Furthermore, the fact that sensory properties *are* basic accounts for why they are unsuited for the role that these fringe experiences play, i.e. to “inflect” or “color” the sensory features that are (among) the direct, first-order objects presented to phenomenal awareness (the “concrete, singular material” of the “sensory manifold,” to quote Joe Levine (2011) paraphrasing Kant). A phenomenally basic sensory property is too fundamental to conscious experience to *merely* modify how we experience another sensory property, without, as it were, contributing something of its own to the phenomenal foundations. So for this task, we need the fringe.⁹⁴

consciousness, we necessarily would be “confused” by the experience, at a loss to know who or what the experience was about.

⁹³ The conservative could claim that the FoF is a type of *emotional* feeling, and argue that these reduce to the sensory (Prinz 2004). This would not alter the fact that identifying fringe feelings with sensory experience is introspectively implausible, for the reasons stated. But also, it is far from obvious that emotional phenomenology *does* reduce to sensory phenomenology. Perhaps instead, fringe feelings and (at least some) emotional feelings form part of a larger family of non-sensory experiences, which fall under the umbrella category of *affective* experience. For a stimulating discussion of affect experience congenial to the neo-Jamesian fringe view, see Jackendoff 1987, chapter 15.

⁹⁴ This characterization of the difference between fringe and sensory phenomenology fits nicely with what Jackendoff (1987) says about affect experience. He is well known for advocating the view that only intermediate-level perceptual representations make it into conscious awareness, but he goes on to

A second reason for thinking that fringe experiences are phenomenally non-sensory comes from research in cognitive psychology and neuroscience. Russell Epstein (2000, 2004) has made a speculative but compelling argument that the neural correlates of fringe experiences probably consist of two principle areas: (a) the hippocampus and related areas in the medial temporal lobes, which may have evolved from a mechanism for navigating through physical space into a general purpose “association machine” that navigates through a network of associated representations stored in long-term memory; (b) mechanisms in the frontal lobes that monitor and control this “mental navigation” process.⁹⁵

If this proposal is on the right track, can it tell us anything about the phenomenology of fringe experiences? Phenomenology that depends on activity in (a) might include fringe feelings that make us aware of other possible directions of thought, locating the current nucleus in relation to various associated unconscious representations, what I earlier called “locating fringe phenomenology”. Given the hypothesis that these cognitive mechanisms use some of the same machinery as is used for external navigation, any phenomenology to which they give rise *might* be a species of sensory-imagistic phenomenology, supposing the underlying representations are sensory. However, Epstein argues that the representations

supplement this view with what he calls the “affects,” many of which correspond to fringe feelings we have been discussing. He motivates the need for this amendment to his theory by observing that the essential difference between certain sensory-perceptual experience pairs—e.g. between an image and a percept of the same visual scene, a voluntary and an involuntary bodily motion, or a familiar image that stands for a memory and a novel one that stands for a new idea—cannot be accounted for by any difference in the “form” of their sensory phenomenology (cf. Epstein 2004, 21-22). So we must appeal to affects, which are “associated with individual entities in the phenomenal field,” and characterize “how they feel” to us (305, 304).

⁹⁵ He argues that this story is consistent not just with the well-known theory of the hippocampus as a “cognitive map” (O’Keefe & Nadel 1978), but also with other theories of the function of the hippocampus, including Jeffrey Gray’s prediction and comparator model (1995; Gray & McNaughton 2000) and declarative memory models (e.g. Squire & Zola-Morgan 1991).

involved in the associative network instantiated in the hippocampus are not themselves sensory, since output from sensory areas in the cerebral cortex—e.g. visual processing areas V2, V4, and MT that support mid-level visual representation—travels through other areas of processing before reaching the hippocampus, with the result that the representations “that the hippocampus operates on to determine the next [nuclear] thought” are “multi-modal and partially amodal” (2004, 23).

If he is right on this point, there are several things we could say about locating fringe (L-fringe) experiences. Epstein himself thinks that the association process that takes place in the hippocampal regions—he describes it as the “progression of thought”—remains entirely non-conscious, on the grounds that higher-level perceptual and conceptual representations are never phenomenally conscious, but only intermediate-level, modality-specific perceptual representations make it into the conscious nucleus and serve to express our thoughts there. This happens when the outputs of the hippocampal memory network processing make it back to the cortex, causing a new conscious perceptual representation to form. But in addition to nuclear experiences, he also acknowledges fringe experiences, which he attributes to the frontal lobes (see below).

Alternatively, one could insist that L-fringe experiences, like nuclear experiences, must be subserved by mid-level perceptual representations realized in sensory areas of the cerebral cortex, in which case, presumably, their phenomenal character is sensory-imagistic.⁹⁶ But this option would undermine the whole point of explaining the experience

⁹⁶ To clear up a possible confusion: in chapter 1, we equated sensory phenomenology with perceptual phenomenology that represents only “low-level” content, e.g. colors, shapes, sounds, etc. But it is plausible that this “low-level” phenomenology corresponds to what in the current context we are referring to as “intermediate-level” or “mid-level” perceptual representations, since these representations seem to best

of thinking, and in particular the introspective distinction between nucleus and fringe, in terms of the different mechanisms of Epstein's neuro-cognitive model. A third alternative is for L-fringe experiences to originate in the hippocampus after all, in which case they presumably would be supported by representations with non-sensory phenomenology. But Epstein could still be right that the higher-level amodal representations in the hippocampus, operations on which determine the direction of the stream of thought, never make it into the stream themselves. For L-fringe experience could correspond, not to these representations, but to the accessed parts of the association network itself—an option that Epstein does not seem to consider.⁹⁷

However, whether or not associative memory processes in the hippocampus support fringe experiences, that still leaves us with the frontal lobe monitoring mechanisms, i.e. (b) in the initial description above. And any phenomenology from (b) is plausibly *not* sensory. For the frontal lobes are involved in higher-level thought and executive functions, so representations that originate there are likely to be very different from the mid-level modality-specific sensory-perceptual representations that arguably underlie (constitute or instantiate) sensory phenomenal properties. Given their hypothesized cognitive role, areas of the frontal lobes are likely to be the neural correlates of the “evaluative” type of fringe experiences (E-fringe), which include the feeling of Rightness, many other meta-cognitive feelings such as the feeling of knowing, and perhaps

characterize what our perceptual experiences, at least of the “low-level” sensible properties, are actually like (for a quick and accessible defense of the intermediate-level idea, see Prinz 2011).

⁹⁷ In that case, the representations that underlie L-fringe experiences might be in an imagistic, map-like format, which represents in “spatial” terms the internal relations holding between stored pieces of information, but without these experiences being of the sensory-imagistic phenomenal type. Could an experience be imagistic, i.e. have a pictorial structure, without being sensory? This is an important question, but unfortunately it lies beyond the scope of the thesis.

even aesthetic feelings of beauty and order.⁹⁸ Indeed, Epstein and Mangan appeal to experimental data to argue that this is in fact the case.⁹⁹

The upshot is that at least one important type of fringe experience (E-fringe experience) likely depends on representations that are non-sensory. Therefore, this fringe phenomenology probably belongs to a family of phenomenal properties that are qualitatively unlike all sensory phenomenal properties (in an analogous way to how, e.g., all auditory phenomenal properties are qualitatively distinct from all olfactory phenomenal properties).

This admittedly speculative picture does not prove that fringe phenomenology is non-sensory. For instance, one could respond by arguing that the results of these cognitive monitoring processes in the frontal lobes, such as meta-cognitive and aesthetic judgments, can only enter the stream of consciousness in “sensory garb”; or one could argue that the evidence for fringe feelings originating in the frontal lobes gives us a reason to think that not all frontal activity is cognitive, since the fringe must be sensory.¹⁰⁰ But it is unclear why we would adopt these positions unless we were *already* committed to the view that phenomenal consciousness is exhausted by intermediate-level perceptual representations. And I see no good reason why we should be committed to this view, even if we are

⁹⁸ Epstein does not distinguish explicitly between locating and evaluative fringe experiences, but what he does say suggests that he thinks all the various types of fringe feelings—e.g. all those on his list of examples—are likely to originate in regions of the frontal lobes.

⁹⁹ Mangan (2001) mentions evidence that deficits in the feeling of knowing are associated with frontal lobe damage (Janowsky et al. 1989). Epstein (2004) appeals to studies that suggest “lateral prefrontal cortex may be involved in detecting ordered patterns in the world” (Bor, Duncan, and Owen, 2001). In general, there is considerable evidence that meta-cognitive experiences such as FoK’s and ToT’s and the general monitoring and control of memory are related to frontal lobe activity (Metcalf 1994; Shimamura 1995a, b; Dolan & Fletcher 1997; Buckner, Kelley, & Petersen, 1999; Moscovitch & Winocur 2002; Schnyer et al. 2005; Bor et al. 2006).

¹⁰⁰ Thanks to Tim Bayne for pressing this second point.

convinced that all *perceptual* experience is intermediate-level representation.¹⁰¹ So the above considerations from cognitive neuroscience, taken along with the introspective evidence, give us enough of a reason to hold our ground on this issue in the face of Conservative skepticism. Thus, since we have reasons to believe that fringe phenomenology is both non-sensory and closely associated with thought, we are justified in thinking of it as a kind of Moderate CP.

IV. How the Fringe Helps Explain Contrast Cases

1. Easy Cases

The Moderate CP account that I am defending is faced with the Conservative challenge just discussed, of explaining why the fringe should count as a version of CP. The account also faces a challenge from the Liberal direction, namely to demonstrate that fringe phenomenology is just as good at explaining the evidence for CP as is the standard Strong CP view that embraces both the Qualitative and Modal Theses. To meet this challenge, in this section I want to consider how the view handles a number of phenomenal contrast cases. In fact, I believe every well-known contrast case that can be explained plausibly in terms of Strong CP can be explained at least as well in terms of fringe phenomenology.

Some of these contrast cases fit explanations in terms of the fringe especially well. For instance, we spoke earlier about Siewert's example of his friend asking him about a borrowed book, and Siewert provides several similar examples of "sudden realizations." One of his best-known examples is the experience of suddenly realizing that he has

¹⁰¹ That is to say, even if we are convinced by the arguments put forward by the likes of Jackendoff (1987) and Prinz (2002, 2011). Again, it is notable that Jackendoff himself, in spite of the other commitments of his theory, allows for a type of experience that is distinct from sensory-perceptual experience, and which bears a marked resemblance to the fringe.

forgotten his briefcase at home (1998, 287ff). In such a case, the entirety of a thought can seem to strike us in an instant. But, I would suggest, what is clear from introspection is only that *something* phenomenal strikes us in an instant, and that then we immediately are able to describe, to ourselves or to others, the entirety of this thought. But it is not obvious that the phenomenal something is an experience of the thought's content. It may be that the grasping of the thought happens non-phenomenally, and the feeling of sudden realization is meant to signal this sudden accessibility of a new thought to the conscious subject. If so, then fringe experiences are ideal for accounting for this rush of feeling.

In the case of remembering the borrowed book, I have already suggested that the feeling of suddenly remembering could be the cognitive feeling of Rightness, that a proposed answer to the unspoken question, "what book?" (and the imagery that accompanies this answer), strikes the subject as obviously correct and appropriate. The feeling of Rightness might precede the imagery of the book by a split second, thus creating the impression of an "ah ha moment", that the subject has "got it", even before details about the right answer have flowed into consciousness. The distinction between the ah ha feeling and the imagery of the right answer is dramatized in situations where one's experience of epiphany is suddenly interrupted by an external distraction, before one has had the chance to articulate the right answer to oneself, but after one has had the feeling. One then feels (often palpably) certain that one had the answer, but must concentrate to consciously grasp it again.

In the case of remembering the forgotten briefcase, however, things might be slightly different, since in this case the realization comes unbidden, and does not follow a conscious memory search. In this sort of case, the most prominent fringe feeling is not

Rightness, but some similar feeling, such as a feeling of salience, that precedes and heralds some important new information that is about to rush into consciousness.¹⁰² The feeling of Rightness might follow at that point, surrounding any sensory images, giving them a “stamp of authenticity,” so to speak. For example, that memory image of your briefcase by the breakfast table is experienced as being obviously accurate and relevant to the present context, i.e. to the experience of being stopped in one’s tracks by a realization.

The puzzle of Thoughtful and Thoughtless, which Nes uses in order to highlight the “thematic unity” of our conscious episodes, is also relatively straightforward to explain. Recall that Nes argues that in order for our sensory and emotional experiences to feel unified around a common topic, the content of that topic must be a constitutive part of the overall experience. I suggested above that he overlooks the possibility that our experiences seem thematically unified by some theme or other, without the details of that theme being consciously accessible beyond the sensory and emotional features in the conscious stream themselves. Now, our sensory imagery in the nucleus can be experienced as unified because of fringe feelings relating them each to a single earlier experience. For instance, Thoughtful’s inner speech, image of the plane wreckage, and feelings of fear are all experienced as *responses* to (his experience of) his friend’s spoken words: that a BA plane just crashed at Heathrow.

Remember, it will not do to object here that, as described, Thoughtful would not be aware of what he is responding to, since his responses are directed merely at spoken words, rather than at their meanings. For we can allow that Thoughtful *knows* what his

¹⁰² The relevant cognitive feeling can be something like Imminence, especially if the feeling precedes the realization by a significant amount of time, “a feeling of urgency inappropriate for the current activity” (Galin 2004, 247). It acts as a nagging request for the subject to shift her attention to discover the identity of this new important information (e.g. the fact of the forgotten briefcase).

friend's words mean without the meaning itself entering phenomenal consciousness. And he experiences the words as meaningful due to cognitive feelings such as Rightness that inflect his sensory experience of the words with the impression that he grasps them. So he does not experience his sensory reactions as occurring simply because a string of meaningless (though syntactically well-formed) noises were emitted from his friend's mouth.

Further, his experiences on hearing the news will feel directly connected to each other in various ways, as when the words "plane crash," spoken over and over in Thoughtful's mind, are experienced as connected by a feeling of Imminence to an as yet unconscious thought—that he might crash on his next flight—that is experienced as pressing in on his conscious awareness. And this sense of Imminence of the foreboding unconscious thought is experienced as negative, and as prompting a surge of anxiety within Thoughtful. When the dreaded thought does enter his conscious stream in the form of visual images of burning fuselage, he experiences it as the object of, as "answering to", his earlier anxiety feelings, and thus as following coherently on from them. So items in our conscious stream are experienced as unified thematically because they are experienced, via the fringe, as directed towards one another, as prompted by one another (or by a common antecedent experience), and as appropriate responses to one another.

In a similar way, the fringe can help explain Horgan's "agentive-cognitive phenomenology." These experiences involve a sense of the overall coherence and purposefulness of our percepts, imagery, and (experienced) bodily movements over time. We can account for these feelings in terms of fringe experiences relating distinct experiences in the nucleus. For instance, when my wife asks me to go get a carton of milk, I

experience my subsequent movement from my armchair to the wardrobe as motivated by her request, because my experience of my own actions is situated between a fringe feeling that “reaches back” to the string of sensory experiences just previous and another fringe feeling that “reaches forward” to various possibilities that seem to be opened up by (my experience of) walking to the wardrobe—even when neither the former sensory experiences, nor the potential future ones, are presently in the nucleus of consciousness. The latter experience is none other than the sense of Imminence we discussed in the last section, while the former experience is a bit like the opposite and complement of Imminence, which we might call “Provenance”: a sense that the current nucleus is not spontaneous or self-generated, but has a cause or origin in something previously in consciousness, which is now only dimly felt in the form of this cognitive feeling. Between them, these two fringe experiences—along with others—imbue the experience of my walking to the wardrobe with a sense of purpose.

Another kind of phenomenal contrast involves the experience of understanding language, such as in Strawson’s example of hearing a foreign language you can speak versus hearing one you cannot speak, and Siewert’s example of reading a difficult passage with comprehension after reading it “spaced out”.¹⁰³ We saw in chapter 3 how the experience of reading with understanding is likely to generate significantly more imagery than the experience of reading without understanding, so Conservatives like Prinz can point to these sensory differences to explain the phenomenal contrast. But I also suggested that if Prinz were right that this is the whole story, then the phenomenal contrast in

¹⁰³ Similar cases involve experiences of lexical or syntactic ambiguities (e.g. garden path sentences, hearing the word ‘virus’ as being about a computer versus hearing it as being about an illness). I think fringe phenomenology can easily be accommodated as part of any plausible explanation of what it is like for the subject in these cases.

question—which, it would seem from experience, is present any time we go back and read with understanding a passage we initially failed to grasp—would *not* be present in many experiences, namely those that do not involve any sensory imagery except inner speech (and perhaps visual imagery of some of the words on the page). For such verbal imagery can equally be present prior to the subject rereading and coming to grasp the passage’s meaning. So Prinz’s sensory-based explanation of these kinds of contrast cases is incomplete.

Here, again, the fringe provides a ready explanation. It may be that the *nucleus*—those features of consciousness that are typically at the center of our attention and provide details of the contents of our thoughts—is populated entirely by verbal imagery in a typical experience of reading a passage with understanding. But even so, the fringe surrounds and suffuses these nuclear elements with cognitive feelings that provide them with a sense of “significance and value.” And, by contrast, we should expect the fringe to be seriously impoverished while we are reading a passage *without* understanding it. The mere visual or aural experience of a group of words—even properly parsed, fluently processed, and read with a natural cadence and intonation—is different than the experience of those words as meaningful, because the latter involves fringe feelings that give the experience of the word the sense of being charged with latent information, with the possibility of the mind accessing additional sensory experiences from memory. This seems to capture what the experience is like when we introspect on it: a word might have a particularly idiosyncratic connotation for us, of which we cannot possibly be aware when we read a sentence that contains the word without understanding it, but when we suddenly grasp the meaning of the sentence, it is precisely this background sense of the connotations of the words that we

gain. The fringe does not make the specifics of each word's connotation conscious, but rather creates the general impression that there is a kind of "halo" of perceived nuance surrounding (at least many of) the words on the page.¹⁰⁴

Finally, what of the last main type of phenomenal contrast we investigated in the previous chapter, what I called Perceptual-Recognitional (PR) contrasts? Two examples we considered were Siegel's story about someone who becomes an expert in recognizing pine trees (a contrast between seeing the trees as a novice, and seeing the trees as an expert who recognizes them as pine trees) and the shift between seeing the duck-rabbit picture as a duck and then as a rabbit. The basic idea was that if we can argue that these cases illustrate an addition (or change) in the "higher-level" (e.g. kind) properties represented by the subject's visual experience, and that representing these properties requires the exercise of conceptual thought, then these cases could be used as evidence for the truth of the Modal Thesis and the existence of Strong CP.

The gist of my response to this strategy was that these cases can be explained either entirely in Conservative terms, by an appeal to shifts in perceptual attention and therefore a difference in sensory phenomenology, *or* in terms of the experience of a "match" between perceptual inputs and non-conceptual representations stored in memory (i.e. what I called "Complex-Sparse-Gestalt" (CSG) representations), which makes a phenomenal difference,

¹⁰⁴ I do not think we must say that the "halo" is entirely general, that it is exactly the same for every word we read with comprehension. There is a lot of space between saying this, on the one hand, and claiming that there is a unique cognitive phenomenal experience, a unique halo, corresponding to *every word* that we read and experience in the sensory nucleus (or every word that represents a unique concept), on the other hand. We can allow that there are all sorts of differences in the fringe while we move from word to word on the page, and that some of these fringe differences correspond to differences in the connotations we associate with particular words. To give two simplistic examples: one word might have a positive connotation for us while another has a negative one; or one word might bear very rich significance for us, due to a long personal history with it, while another word might bring with it few associations for us beyond a rudimentary dictionary definition.

but not one that requires the addition of higher-level content. The latter response, to the degree that it acknowledges a change in phenomenology that goes beyond a shift in which low-level perceptual properties are being attended to, might be thought of as invoking a species of Moderate CP. Explaining the purported phenomenal change in terms of fringe phenomenology is one way to cash out this view.

The most obvious way that the phenomenology of “matching” a CSG representation to a feature of one’s current perceptual experience could be thought to involve fringe phenomenology would be for the experience of recognizing a familiar kind of object to involve a cognitive feeling of recognition (FoR). In this case, the feeling of recognition is experienced as infusing the shape and features of the tree, or the lines of the rabbit drawing, such that we do not just feel an unattached sense of recognition, but we experience the sensory features themselves as familiar.

More precisely, we experience the remembered characteristic(s)—i.e. the complex-sparse-gestalt property, retrieved from memory—in the visual experience. For the CSG property is composed out of sensory features that are general enough to be consistent with the specific set of sensory features that come from the subject’s current visual input, and the two sets of visual features are experienced as “layered” one on top of the other, as the CSG property is “superimposed” onto the visual experience of the tree. We cannot really separate them in introspection.¹⁰⁵ Yet, because the CSG property is imbued with the fringe feeling of recognition, we are aware of both the particular pine tree we see before us, *and* the remembered tree (or the complex generic property of multiple pine trees past) in the same overall conscious experience. In other words, the nuclear experience combines a

¹⁰⁵ Cf. Briscoe’s (2008, 2011) discussion of “make-perceive”: hybrid experiences that combine perceptual and imagistic elements.

percept and an image from memory. In this way we are able to experience the pine tree before us as being the same—in the sense of sharing the particular properties to which we are attending and on which the FoR is directed—as that which we remember.¹⁰⁶

2. A Harder Case

I think a more challenging kind of case, also due to Charles Siewert, is what we might call “condensed thoughts.” For example, Siewert recounts the following experience of entertaining a complex thought in the period of time it took him to walk from his restaurant table to the cashier:

Asked to state more precisely what this [thought] was, I would have to say something like: “My preoccupation with the topic of my book has made the world seem especially alive with examples of it, references to it, so that it can’t help but seem to me that the world is more populated with things relevant to it than previously. And it struck me that this is similar to the way in which new parenthood made the world seem to me burgeoning with babies, parents, the paraphernalia of infancy, and talk and pictures of these.” Somehow this thought of my philosophical preoccupations and parenthood, and an analogy between their effects, rather complex to articulate, occurred in a couple of moments while I approached the cashier, in the absence of any utterance... [T]hought occurs, wordlessly, without imagery, condensed, and evanescent. If you agree that you have such un verbalized noniconic thoughts, and the way it seems to you to have them differs from the way it seems to have imagery and sensory experience, then you will agree that noniconic thinking has a phenomenal character distinct from that proper to iconic thinking and perception. (Siewert 1998, 277-278)

I suspect we all know the sort of experience that Siewert has so perceptively and vividly described in this passage, and I believe such experiences provide one of the most compelling reasons to accept the existence of *sui generis* phenomenology of thought. The case is arguably more difficult to explain without appeal to Strong CP than any of those we

¹⁰⁶ The feeling of recognition seems akin to the feeling of familiarity we experience when we see a face that looks familiar but cannot place it, except that in the case of the pine tree and duck-rabbit, we *can* place what we see.

have considered already, because as described it is hard to deny that it's the thought's *content* itself that is the dominant characteristic of the experience.¹⁰⁷

How exactly should we account for condensed thoughts? First of all, we should probably treat Siewert's contention that condensed thoughts occur without any imagery with a dose of skepticism. That we sometimes entertain complex thoughts without articulating them in complete sentences in inner speech is one thing; that we entertain such thoughts without *any* imagery is much less plausible, and Siewert is not entitled to present this detail as common ground in his description of the case. Second, the fact that we can have conscious awareness—in the sense of *access* consciousness—of the full, complex content of a thought, like the one Siewert describes, in a matter of moments, does not show that we have an *experience* whose phenomenal character entails we are representing the entirety of this thought content. The striking experience Siewert describes, of grasping a thought in a flash in all its detail and complex structure, might *accompany* the non-phenomenal representation of the thought content, without itself entailing, or being constituted by, such representation.

Nevertheless, there does seem to be something in the experience of a condensed thought that the paucity of information given by sensory imagery cannot capture. After all, the experience Siewert describes is one of being *struck by a thought*, which on the face of it is very different from the momentary enjoyment of a few fleeting images flashing through the mind (e.g. a visual image of a baby's face, verbal imagery of a sentence fragment, etc).

¹⁰⁷ The observant reader will notice that this example isn't exactly a contrast case, but more of a bald assertion about what it is like to think certain thoughts, presumably in the hope that others will find the description compelling based on their own experiences. Thus, when evaluating the example as an argument for CP, it's hard to say at the outset what should count as a successful explanation of the case, since it's hard to say what needs to be explained without begging the question one way or the other (cf. Robinson 2005, 553ff).

This difference, between the experience of the activity of thinking and the experience of a mere string of images, is made especially vivid in a case like this in which the thought is especially complex and the period of time it takes to consciously entertain it is especially brief. For it is the swiftness of the thought, that we grasp the whole of it in little more than the blink of an eye, that makes its complexity stand out. And it is the complexity of the thought, the sense that what has just occurred to one contains various moving parts, various rebuttals and analogies and sub-clauses, that makes it seem like one's experience of this thought could not possibly be exhausted by a few images. In other words, I think what is striking about the kind of experience that Siewert's example highlights is that the overall experience of grasping the thought involves an awareness of the "architecture" of the thought as a whole.

Fringe phenomenology can account for this sort of experience of the structural elements of a thought. As we noted earlier, the experience of relations and transitions between thoughts, and of the anticipation and awareness of distinct parts of a complex thought, is a fringe phenomenon. Indeed, something like it is mentioned several times in Epstein's list of prominent fringe feelings reproduced above (5-7 on the list). For instance, Epstein talks about "the sense we have for the overall scheme or form of a book, a work of art, or a scientific system." While I am not sure I consciously experience the "overall scheme" of a scientific system, I do think I can become conscious of how an overall argument hangs together, and that this is what Siewert experienced when he was struck by the analogy between work on his book and new parenthood. I do not mean that anything as detailed as the complete and precise logical structure of a thought is made manifest to the subject in a fringe experience. But *some* structural elements of the thought can be

directly felt: that the first part of my thought is alike in certain respects to the second part of my thought (perhaps because the relation that holds between elements in the first part is the same as the relation that holds in the second part), that point X is a good rejoinder to anticipated objection Y, and so on. In each case, the experience is made up of fringe feelings, since what is felt are not the contents of each part of the thought—they show up as images in the nucleus, if at all—but the relations between the parts. Locating fringe feelings, like Imminence, can help to represent these relations, by representing connections between (past, present, and future) sensory images.

In summary, the fringe account of this case claims (a) Siewert had an experience of the thought in schematic form with L-fringe phenomenology combined with imagery. It is notable that Siewert's thought has the form of an analogy, which would lend itself especially well to a diagrammatic representation of relations between locations in mental "space." And further, (b) this experience was combined with various fringe feelings of the meaningfulness and fittingness that comes with having articulated to oneself an especially interesting idea, particularly one that captures a certain order or pattern observed in the world (in this case, an analogy between two superficially disparate domains).¹⁰⁸

Siewert no doubt would object that this account leaves out his awareness of the *particularity* of the thought itself. But his impression that he was consciously experiencing the particular details of the thought can be explained by a feeling that he was, in that moment, "tracing" the distinctive contours of a particular unconscious content. The *outline* of the thought was present in the phenomenal stream—in the way described in the previous few paragraphs—but along with it there was a palpable sense that what was

¹⁰⁸ Cf. Jackendoff's description of the affect he calls "*congruous*" (1987, 308-309).

presented in conscious experience was *answerable to* something more, something real and present, though still “under the surface”: the detailed thought content that was immediately to hand, and which he immediately could have gone on to articulate in inner or outer speech.¹⁰⁹ And in fact, this feeling was veridical: the thought *was* directly at his control and disposal, even as the details of its contents remained largely phenomenally unconscious.

Indeed, I think the fringe explanation is more plausible than Siewert’s alternative, since I strongly suspect that he could *not* have entertained in consciousness the entirety of his complex thought, which it takes him a paragraph to describe in words, in the two to three seconds it took for him to walk to the restaurant cashier. I will present my reasons for thinking this is impossible in the next chapter. If his response is that the experience of the thought is somehow “condensed”, then I would point out that how a condensed experience nonetheless is able to individuate a *non-condensed* thought—i.e. to make the full content of this thought manifest to the subject in consciousness, in such a way that the experience determines that it is *this* content that is thought, rather than any other—is left entirely mysterious on the Strong CP account. So a plausible middle ground is to hold that there is more than imagery to condensed thoughts, but less than Siewert imagines.

¹⁰⁹ Farid Masrour (2011, 12-13) provides an interesting example which illustrates this sense of having a thought “to hand” prior to it entering conscious awareness: suppose one is listening to a philosophy talk and in a flash intuitively objects to the speaker’s argument, but as one is still working to articulate the objection in inner speech, someone else from the audience raises their hand and gives the same objection. At that moment, one can have the distinct sense that the other person’s words exactly capture the thought that one had (alternatively, that her thought is similar but different to the thought one had) even though one never succeeded in fully articulating the thought to oneself. In that instant, one’s feeling is a fringe experience of Rightness or fittingness, expressing that the other person’s words “click into place” with one’s thought. This feeling suggests there is something outside of consciousness to which one is comparing her statement, to which her words are “answerable”.

This completes our survey of phenomenal contrasts that are commonly used to argue for cognitive phenomenology, and our consideration of how fringe phenomenology can explain these cases.

V. Conclusion

In this chapter I have sought to draw attention to an oft-neglected type of phenomenology and its relevance to the CP debate. Given such arresting examples as FoF's, ToT's, and déjà vu, it is extremely difficult to deny that at least some of the phenomena described by James and others as "the fringe" are real experiences with genuine phenomenal character. But beyond this fact, I have proposed that fringe experience is far more ubiquitous and important to conscious cognition than some of these more conspicuous and peculiar instances of it might suggest. I have argued, against the Conservative, that fringe phenomenology is genuinely non-sensory and plays a key role in our conscious monitoring and control of our thought processes, and so deserves to count as a Moderate phenomenology *of* thought. And I have argued, against the Liberal, that this type of Moderate CP offers an alternative to Strong CP as a plausible explanation for any of the prominent phenomenal contrast cases used to support the Liberal's position.

I have not been concerned to prove that this position is *superior* to the standard Liberal and Conservative positions. I have been content to show that it is a viable option that deserves attention (more attention than it has heretofore received). However, in the final chapter of the thesis, I present an argument against the Modal Thesis and Strong CP, which I believe counts decisively against that view. If I am right, then Phenomenal Liberals

have a very good reason to take this neo-Jamesian Moderate CP position a lot more seriously.

Chapter 5

Why Phenomenology Cannot Determine the Content of Thought

I. Introduction

1. Up to this point in the thesis, I have offered considerations that suggest there may be phenomenal properties that count as CP according to the Qualitative Thesis, but I have argued that nothing in the current literature clearly supports the Modal Thesis. While evidence for the former thesis—particularly relating to “the Phenomenal Fringe”—is interesting and promising, it is the lack of evidence for the latter thesis that I want to emphasize. In the context of my overall goal to build a case against the Modal Thesis, the preceding chapters can be seen as “playing defense”: demonstrating the inadequacy of current Liberal arguments to support a Strong CP conclusion. In this chapter I go on the offensive, presenting an original positive argument for the conclusion that the Modal Thesis is false: there is no such thing as phenomenology that metaphysically determines thought.

Here is a brief outline of the intuitive reasoning of my argument, to give the reader a flavor of how it will go before we dive in. The argument starts from the observation that the Modal Thesis depends for its plausibility on a claim about the *Phenomenal Intentionality* (PI) of thought. Phenomenal Intentionality is the claim that (a type of) intentional content is either identical to, or somehow metaphysically depends on, (a type of) phenomenology. As the Presentation thesis introduced and defended in Chapter 1 states, all phenomenology, by its very nature, is exhausted by how things seem or appear to the subject. The phenomenal, or qualitative, features of an experience are presented to the subject, and thus are features of what the subject is directly aware of. So in order for PI with respect to *thought* to be true, such that phenomenal features alone make up (or “ground”) the conceptual content of a subject’s thoughts, *all* of the various aspects of that content—or at least enough detail to individuate the content and therefore distinguish it from *every other possible* conceptual content—must be “presented” to the subject in the way discussed, i.e.

they must somehow be included within the subject's phenomenal awareness. In short, what is *thought* must be wholly captured in what is *experienced*—that of which we are aware in the special, phenomenal sense. But, as I will argue, this view of conceptual content is wildly implausible (in a way that the analogous view of perceptual content is not). Thus, we have no choice but to reject PI when applied to thought. And since the Modal Thesis is entirely unmotivated and *prima facie* implausible without the corresponding PI thesis, we should reject it as well.

2. Cognitive Phenomenology and Phenomenal Intentionality

In this sub-section, we will begin with the Modal Thesis, and then show how a PI thesis motivates and supports it. Here again is the Modal Thesis as we stated it previously:

Modal Thesis: there exist phenomenal properties that are *necessarily* unique to, and inseparable from, thought: the instantiation of one of these phenomenal properties in a subject is sufficient (in every possible world) to determine the conscious representation of a thought content, i.e. the tokening of a conscious thought in the subject

According to the Modal Thesis there is a type of phenomenal property such that, in every possible world, if a subject is undergoing an experience with that phenomenal property, that subject is also simultaneously having a thought. That is *all* that the Modal Thesis claims. Thus, the relationship between phenomenology and thought described in this formulation of the thesis is intended to be neutral on questions of why the type of experience is always found together with a thought in a subject, and on the question of the exact relationship that holds between the phenomenal character of the experience and the content of the thought.

The Modal Thesis does a good job of capturing the notion that cognitive phenomenology is proprietary to the cognitive domain, since, while it allows that one may have instances of thought without CP experience, it entails that one will never find CP independent of thought. So the former can justifiably be said to “belong to” the latter. But this fact itself is left entirely unexplained by the thesis. An explanation along these lines is desirable for the Phenomenal Liberal, for at least two reasons. First, because without one the link between CP and thoughts would be just a brute metaphysical necessity and thus highly questionable: why *couldn't* a conscious infant or animal, for example, instantiate these particular phenomenal properties despite not being capable of conceptual thought? Second, an explanation is desirable, because without one CP would lose much of its potential explanatory power. Many supporters of CP hope it can provide the key to a theory of the foundations of thought itself.¹¹⁰ But crucial to such a theory is the idea that CP can account for how thought states are able to possess content in the first place. If we give up on this idea then it is hard to see how Strong CP really helps much in our quest to understand the foundations of thought. We would know that certain experiences are only found with thought (though not necessarily vice versa), but the relationship between the two would be left entirely un-illuminated. How can we move from the Modal Thesis to a claim that would promise to offer a degree of illumination?

First, note that it is difficult to see how instantiating CP properties could ensure that thought occurs, i.e. that thought content is entertained in general, without necessitating that *specific* thought contents are entertained. For how would this work? Could there be a kind of generic phenomenal “buzz” that makes it so that a subject has to be thinking, but

¹¹⁰ For instance, see Horgan and Tienson 2002; Kriegel 2013; Pitt 2004, 2009; Strawson 2011.

not thinking about anything in particular? By stipulation, the buzz could not accomplish this feat by making the subject entertain any specific content (and what would it be to entertain a wholly non-specific content?), but if it couldn't ensure that some content was entertained, then how could it necessitate that any thought was going on at all? Of course, the buzz could necessitate that *some specific content(s) or other* be entertained, without fixing the specifics all by itself. But then what factors *would* fix the specifics? Whatever additional factors we posit to help fix the content, we still will be in a position to wonder why the CP buzz could not occur in the absence of those factors, in which case we still do not have a good reason to believe that the presence of CP alone suffices to guarantee the entertaining of thought contents.

For instance, one might propose that the buzz could be a kind of *attitude* experience, where the specifics of the contents to which the subject bears the attitude are determined by something other than the buzz. But why suppose such attitude phenomenology couldn't be instantiated (perhaps only in some bizarre metaphysically possible scenario) in the absence of an occurrence of genuine thought? The subject could have an experience as if she were entertaining some quite determinate content or other, but be mistaken. Her mind would be "empty," with no awareness of a meaning for the attitude experience to latch on to. It is hard to see how such a scenario could be ruled out from the outset, if all CP amounted to was a quite general attitude phenomenal property (or several more specific attitude properties).

The Phenomenal Liberal can address this issue by adopting a Modal Thesis about the relationship between specific CP experiences and specific thoughts, something like the following:

Specific Modal Thesis: there exists a class of phenomenal properties that are necessarily accompanied by specific cognitive states: for each phenomenal property type, P, within this class, there is a distinct conceptual content, C, such that the tokening of P in a subject, S, necessarily determines the tokening of a conscious thought in S that has C as content.

If this version of the thesis is true, then the original Modal Thesis must be true as well. But we still do not have an explanation for why CP experiences determine that specific conceptual contents are represented in the first place. For that, we can appeal to a *Phenomenal Intentionality* (PI) thesis regarding thought. It is noteworthy that most of the proponents of CP are also defenders of some form of PI view. What all PI theses share is the claim that there is a kind of intentionality, so-called *phenomenal intentionality*, that is inseparable from, intimately connected to, (a kind of) phenomenal character. There is little consensus about the best way to characterize this intimate connection, but for simplicity's sake, I will adopt the following PI thesis for thought:

Cognitive Phenomenal Intentionality (C-PI): an experience with a particular cognitive phenomenal character—i.e. the instantiation of a particular CP property—*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 (or impure representationalism) is the thesis (for a class of phenomenal properties) that all phenomenal properties (within that 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, where 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.¹¹¹ In the case of C-PI, the particular 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. This is not to deny that there are wide thought contents as well. In general, we can leave open the possibility that combining the CP facts with certain other facts about the subject or his environment—e.g. non-cognitive phenomenal facts, or causal-functional facts—might fix one or more layers of conceptual content that extend beyond basic cognitive-phenomenal content.¹¹² But the important thing with respect to C-PI is just that there is *a* genuine type of intentional content the representation of which is *wholly* a matter of the instantiation of cognitive phenomenal properties in the subject's mind.¹¹³

Must the phenomenal and intentional properties in question be strictly speaking *identical*? We said that the motivation for adopting C-PI is to have a more illuminating account of why the Modal Thesis is true, i.e. why CP properties determine thoughts. We *could* respond to the need to explain the Modal Thesis by construing the C-PI thesis as the claim that CP properties (metaphysically and explanatorily) “ground” intentional properties, rather than being identical to them. I cannot see any real harm in opting for this alternative, and in fact, much of what I will go on to argue in this chapter could (with some inconvenience) be restated in terms of grounding or some sort of dependence relation,

¹¹¹ I borrow the notion of manner from Chalmers 2004, and mode from Crane 2003 and 2014, chapter 8.

¹¹² Horgan & Tienson (2002) and Siewert (2011), among others, discuss the relationship between CP contents and wide thought contents.

¹¹³ Note too, as we discussed in Chapter 1, the cognitive phenomenology in question could be *impure*: it may not be possible to separate out in experience the CP properties from other, non-cognitive phenomenal elements (Levine 2011).

rather than identity (or in terms that are neutral between the two). But without a more detailed description of that aspect of the CP character's intrinsic nature that somehow confers upon it the power to necessitate the representation of thought contents, the grounding claim would still leave the reasons for the truth of the Modal Thesis opaque. By contrast, the identity claim makes the reasons for the inseparability of CP and thought content entirely transparent. Of course, there is never really an explanation for why one thing is identical to another thing. But *if* a thought's phenomenal properties are identical to its (impure) intentional properties, then clearly for a subject to undergo a CP experience she would also have to think a thought.

Additionally, the idea that CP properties are identical to *impure* intentional properties, rather than pure intentional properties, makes the identity claim less radical than one might suppose, since a lot can be built into the requisite manner of representation, so that not just any representing of a thought's content suffices for a cognitive phenomenal experience. Once we build in this reference to the distinctive cognitive-phenomenal *way* that a CP experience enables a subject to be intentionally directed onto a thought content, I see no harm (at least for the advocate of PI concerning thought) in conceiving of CP properties as strictly identical to intentional properties of a certain kind.

One might worry that construing C-PI as an identity claim does not allow for the phenomenal properties to take explanatory or metaphysical *priority* over the intentional properties, the way a grounding claim would, since unlike identity, grounding and other dependence relations are asymmetric. This could be seen as a problem, since Phenomenal Intentionality is often described as a view on which phenomenology is somehow more

basic than intentionality. But in fact, proponents of PI theories are often unclear or uncertain about whether the view should involve thinking of phenomenology as a kind of intentionality or just as somehow being a basis for it.¹¹⁴ What *is* clear is that PI theorists hold that understanding phenomenology should be a primary starting point in the philosophical and scientific study of the nature and origin of intentionality. This general approach has been dubbed by Uriah Kriegel (2013) the “Phenomenal Intentionality Research Program,” and is nicely summed up by Adam Pautz’s (2013) memorable slogan, “Phenomenology First.”

In light of this basic outlook, I suspect that what is essential for *many* PI theorists is just that, in spite of the fact that intentionality is somehow inextricably wrapped up with phenomenology, the latter cannot be explanatorily reduced without remainder to representation described in entirely non-phenomenal terms. And our formulation of C-PI meets this requirement, since it is not just a form of impure intentionalism but a form of *non-reductive* intentionalism as well: our specification of the manner of representation for CP experiences makes essential reference to phenomenology, so it does not allow for an account of CP entirely in non-phenomenal intentional terms.¹¹⁵

Furthermore, many proponents of PI take phenomenal intentionality to be the source of all other, merely derivative, kinds of intentionality, for example the intentionality of natural languages and human artifacts. If we combine this point with the fact that on the PI view phenomenology is not reducible to intentionality, we can see why PI proponents

¹¹⁴ For instance, Kriegel 2013 states PI as a grounding thesis, but then allows that on this use of the term, grounding does not rule out the possibility that the intentional property grounded by a given phenomenal property is identical to it. Chudnoff 2015 compares and contrasts PI grounding, identity, and necessitation theses.

¹¹⁵ Cf. Byrne 2013.

take phenomenology to be central to the correct account of how intentionality first gets “injected into the world” (Kriegel 2013). The explanatory usefulness of this story comes not in offering a reduction or analysis of intentional phenomena in purely non-intentional terms, but in *locating* the origins of the intentional in the phenomenal-intentional. In this way, the PI view is importantly different from typical “naturalizing” accounts of intentionality, which seek to reduce it to the presence of causal or other kinds of natural “tracking” relations that hold between mental states and features of the subject’s physical environment. Thus, in this sense at least, “Phenomenology First” is an apt slogan for all Phenomenal Intentionality views, even those, like C-PI, according to which neither the phenomenology nor the intentionality is more basic than the other.

In short, the C-PI thesis serves to explain why the Modal Thesis is true (if it is), something that otherwise is not easy to do. Although one might quibble with the details of how C-PI is stated, I think almost any Phenomenal Liberal who believes in Strong CP will accept a thesis similar to C-PI. For if we reject C-PI and related theses, the Modal Thesis on its own becomes a much less attractive claim to defend. Indeed, it is unclear why one would be inclined to believe it. Thus, if in this chapter I can demonstrate that C-PI is untenable, the Modal Thesis will become virtually untenable as well.

3. The Standard Account of C-PI

Before moving on to my main argument, I need to flesh out the view of the relationship between CP and thought that naturally follows from the Cognitive Phenomenal Intentionality (C-PI) thesis. I want to explore a few of the implications of C-PI that I think

cause trouble for the position. Just to have a label, we can call the picture of C-PI that emerges the Standard Account (SA) of Cognitive Phenomenal Intentionality.

First, representing a thought content in an essentially phenomenal and cognitive manner is plausibly the same thing as consciously *entertaining* a thought, where by “entertaining” I mean to pick out a cognitive kind of occurrent awareness of content. Conscious thinking is cognitive awareness, and a thought is a state or episode of such awareness. But now, second, recall that in the first chapter I argued for the following thesis about the nature of phenomenology:

Presentation: all phenomenology is presentational in nature, consisting exclusively in the apparent presentation of objects (i.e. particulars and properties) to the subject 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.

I stressed that accepting the kind of phenomenal awareness described by Presentation does not require us to accept any additional controversial claims about the nature of this awareness, for example that it should be assimilated to a perceptual model of acquaintance, or that the “objects” of awareness are necessarily mind-independent, public, concrete, external entities. Rather, the awareness in question is merely the result of there being something it is like for the subject, in the usual phenomenal sense, to which the subject can directly attend.

So both experiences with phenomenal character and thoughts can be understood as being kinds of awareness. Combining this point 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) a state/episode of being aware of a particular thought content in a distinctively phenomenal and cognitive way.

The Awareness thesis replaces the descriptions of CP experience and the representation of thought content found in the original C-PI thesis with descriptions of the two types of awareness: the phenomenal awareness of the qualitative, what-it's-like character of a CP experience, and the cognitive awareness (i.e. entertainment) of the conceptual content of a thought. But according to Awareness these two types of awareness are identical.

Therefore, assuming Presentation is true and that thinking a thought amounts to a kind of awareness of content, the proponent of Cognitive Phenomenal Intentionality is going to be pressed into accepting Awareness as well. But if Awareness is true, what is the relationship between, on the one hand, the presented qualitative properties that constitute the phenomenal character of a CP experience, and, on the other hand, the properties that constitute the thought content we entertain when we undergo this CP experience? In other words, if a particular state/episode of phenomenal awareness is identical to a particular state/episode of cognitive awareness, what is the relationship between the respective *objects* of awareness in each case? I think the overwhelmingly most plausible answer to this question is some version of the following transparency thesis with respect to CP experience:

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.
- (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.¹¹⁶

¹¹⁶ This thesis is less strict than many transparency theses, which often involve denying that the subject can attend to what I am calling here "manner," i.e. any intrinsic features of the experience itself as opposed to its

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 amount to the same thing as the awareness of what I am thinking about, unless (at least some of) *what* I am aware of is the same in both cases. It may be that, strictly speaking, Transparency is not forced upon us by this consideration. Perhaps one could allow that the awareness of specific cognitive phenomenal character is (type) identical to awareness of some thought content or other, but deny that the objects of phenomenal and cognitive awareness are introspectively indistinguishable, if one thought that awareness of the first set of objects, the presented qualitative properties, somehow acts as a vehicle that conveys the second set of objects, the content properties, to our awareness (or vice versa). The idea would involve accepting (2) in the above formulation of Transparency, but denying (1): a subject is aware of the contents of a thought *through* awareness of the qualitative features of the thought experience. This suggestion strikes me as being of dubious coherence, since it is unclear how we could conclude that there is really only one type of awareness in this case rather than two.¹¹⁷

content. Strong representationalists about phenomenal consciousness have typically claimed that perceptual experience is transparent or "diaphanous" in the sense that a subject, in undergoing an experience, is *only* directly aware of features of the external world that the experience represents (for discussion, see Moore 1903, Harman 1990, Tye 2002, Kind 2003, Siewert 2004, Stoljar 2004, and Pautz 2007). Following Crane (2014, 159-165) and Siewert (2004), I am allowing that the manner or mode, the "particular way of being aware" of a content, can make its way into the phenomenal character of the experience and thus into that of which we can be directly phenomenally aware. But the important point is just that *all* of the thought content makes it into phenomenal awareness as well, i.e. constitutes aspects of what it's like for the subject. In addition, the manner of awareness is still an aspect of the "intentional nature," of the experience, built into the structure of the subject's point of view on the world. It describes *how* one is aware of a content; thus one cannot attend to manner wholly *independently* of attending to one's awareness of content, and thus to *what* one is aware of, as (2) in the formulation of Transparency is meant to underline (cf. Siewert 2004, 35).

¹¹⁷ In response to this worry, one could try reformulating the alternative view as a grounding thesis rather than an identity thesis: rather than Awareness saying phenomenal awareness is identical to cognitive awareness, the claim could be that the first sort of awareness grounds the second.

But in any event, the deeper challenge for this alternative view is to explain how the phenomenal awareness of certain qualitative properties could necessarily mediate the cognitive awareness of certain thought 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 relationship between an arbitrary representation, such as a word or symbol, and what it happens to represent—a merely contingent relationship of indication. 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 of the qualitative properties 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 thinking experience. It would not be sufficient, by contrast, for us to have to make a *judgment* about what we were thinking about on the basis of our awareness of what it is like to think the thought. For in addition to the fact that this idea is psychologically unrealistic, it does not rule out the possibility of a subject making deviant interpretive judgments and thus does not guarantee that awareness of a particular CP qualitative feature would always be transposed into awareness of a particular thought content.

I am not saying that an explanation for how this could work—how the act of “interpreting” the objects of phenomenal awareness as objects of cognitive awareness

could be built into the experience itself—cannot be given.¹¹⁸ But a much simpler and more attractive 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 the experience is like.¹¹⁹

I think it is safe to assume the proponents of Strong CP and C-PI will accept Awareness and Transparency, since several prominent Phenomenal Liberals already have advocated very similar views. For instance, Siewert argues that when we introspect the only “subjectively discernible differences” we find that are “identifiable with the phenomenal differences that occur as we think and understand,” are indistinguishable from “differences in ways of thinking and understanding, where these are characterized in a manner essentially tied to...what is thought/understood (differences in content)” (2011, 27). He concludes that what it is like to think or to understand language is always constituted by what it is like to entertain a specific thought or to take a bit of language to have a specific meaning, which essentially involves awareness of conceptual content.

¹¹⁸ The most straightforward and intuitive way of doing so is something like the following: even though the objects of phenomenal and cognitive awareness are distinct, one can become aware of the latter if one attends to the former and automatically “interprets” them *as if* they are the latter—provided that the qualitative features, and the relations between them, are suitably similar (in important ways that need to be spelled out) to the features of thought content. CP experience essentially involves “taking” the objects of phenomenal awareness to be the objects of one’s thoughts. David Pitt articulates something like this view for visual experience in reply to Michael Tye’s claim, “to suppose that the qualities of which perceivers are directly aware in undergoing ordinary, everyday visual experiences of things are really qualities of the experiences would be to convict such experiences of massive error,” (2002, 139). Pitt responds, “on the contrary, I would argue, if visual experience were not systematically misleading—if it didn’t present itself as something it’s not (namely, external reality)—it would be useless. Transparency is an illusion made necessary by the fact that what experience is supposed to represent is external to the mind, while experience itself is internal. Perceptual experience cannot present itself as what it is if it’s to be a naively believable guide to what it’s not” (2011, 154, fn. 12). But if we were to adopt a view of CP experience that is roughly analogous to Pitt’s view of visual experience, the worries I raise for the Standard Account of C-PI in the second half of the chapter would carry over to create serious trouble for this position as well.

¹¹⁹ This account, and especially the idea that phenomenal character does not have to be supplemented with an interpretation in order to present us with a content, is intended to echo claims Charles Siewert makes about the phenomenal intentionality of perceptual experience (1998, chapter 7).

Similarly, in their influential introduction to cognitive phenomenology and phenomenal intentionality, Horgan and Tienson speak of “the what-it-is-like of experiencing a specific propositional-attitude type vis-a-vis a specific intentional content”:

...if one contrasts wondering whether rabbits have tails with thinking that rabbits have tails, one realizes that there is something *common* phenomenologically, something that remains the same in consciousness when one passes from, say, believing that rabbits have tails to wondering whether rabbits have tails, or vice versa. 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. This aspect of the overall phenomenology of intentionality is *the phenomenology of intentional content* (2002, 522)

They describe a type of phenomenology as “essentially a what-it’s-like of holding before one’s mind a specific *intentional content*,” which seems equivalent to my suggestion that what it is like for us to think is constituted by our awareness of a thought’s content (perhaps together with properties that characterize *how* we are aware of it) (541, note 14). Finally, David Pitt is the most explicit of these authors about holding views akin to Awareness and Transparency: “it is only because type-distinct conscious thoughts have type-distinct phenomenologies (of the cognitive sort) that one can Immediately distinguish them from each other; and it is only because a conscious thought that *p* has a phenomenology that constitutes its (representational) content that one can Immediately identify it as the thought it is” (2004, 7-8). In other words, one’s introspective identification of one’s thoughts amounts to the “Immediate” kind of awareness (he elsewhere calls it a form of “acquaintance”) that one has of phenomenal character, because the content of one’s thoughts just is a distinctive type of phenomenal character, a way that it is like for the subject.

To sum up, what I am calling the Standard Account of Cognitive Phenomenal Intentionality, which has a close affinity to the views of a number of Phenomenal Liberals, including Siewert, Horgan & Tienson, and Pitt, is comprised of the following:

The Standard Account (SA)

- C-PI, and two subsidiary theses:
 - Awareness
 - Transparency

4. Let us review. We started with the Modal Thesis as our primary target, but we said that proponents of this claim would almost certainly only accept it because they also accepted a stronger thesis, namely (some version of) Cognitive Phenomenal Intentionality (C-PI). While the Modal Thesis only posits a bare determination relation between cognitive phenomenology and thought—any instance of a CP experience (of a particular type) requires that there also be an instance of a thought (of a particular type)—the C-PI thesis accounts for this relation by appeal to the fact that the experience and the thought are identical. We then saw that several further interesting claims follow from C-PI (together with plausible assumptions), which taken together make up what I called the Standard Account (SA) of C-PI: first, Awareness, based on the idea that both CP experiences and thoughts can be understood as states/episodes of subjective awareness; second, Transparency, based on the idea that, since these two states of awareness must be one and the same, the objects of awareness in each case, the presented qualitative properties and the represented content properties, should be the same as well. The point to draw from this line of argument is that Phenomenal Liberals who embrace the Modal Thesis are only likely to do so if they also accept the Standard Account, and in fact, I just gave several

examples of Liberals whose stated views do align with SA. But now I want to argue that the Standard Account is very probably false.

I now can sketch the central argument of the chapter, an argument against C-PI, and thus against the Modal Thesis:

Why CP Cannot Determine the Content of Thought:

- P1** The Modal Thesis is probably false, or at least entirely unmotivated, unless C-PI is also true
- P2** If C-PI is true, then, on very plausible assumptions, the Standard Account of C-PI is also true
- P3** But the Standard Account—C-PI, Awareness, and Transparency—is incompatible with all known viable theories of conceptual structure
- C1** Therefore (from P3), SA is probably false
- C2** Therefore (from P1, P2, and C1), the Modal Thesis is probably false, or at least entirely unmotivated

I take it a reasonable moral to draw from C2 is that we should reject the belief in phenomenology that can determine the representation of thought contents. The argument is only intended to be an instance of all-things-considered inference to the best explanation (one *could* object that we simply haven't hit upon the correct theory of conceptual structure yet, for instance). Still, I think the objections against the tenability of the Standard Account (Premise 3) are decisive. I have already explained why I think proponents of the Modal Thesis should be committed to C-PI and SA (Premises 1 and 2). In the remainder of the chapter, I will explain and defend the argument's crucial third

premise, in order to complete my case against the claim that cognitive phenomenology can determine the content of thought.

II. Cognitive Phenomenology and Conceptual Structure

1. Recall from Chapter 1 that we are understanding all mental contents—and the contents of our conscious experiences in particular—as *satisfaction conditions* or *conditions on extension*. Specifically, propositional thoughts, and typical perceptual experiences, have accuracy or truth conditions: their contents, the information they carry and convey, are just the conditions under which they are accurate or true. This conception as I adopt it is neutral as to whether these contents are (complexes of) worldly entities, i.e. the properties, objects, and states-of-affairs to which our words and mental states refer (Russellian contents); or instead should be thought of as modes of presentation of those entities, (Fregean contents); or just as sets of possible worlds; or whether our mental states should be thought of as having multiple kinds of content. I take it each of these views about the nature of contents is consistent with also thinking of them as being, or being given by, conditions on extension.¹²⁰ For convenience in what follows, I will tend to speak of the phenomenal contents of cognitive experiences—those contents that constitute what it is like to think a thought—as being composed of properties, rather than as modes of

¹²⁰ For instance, supposing that modes of presentation should not be understood as definite descriptions or complexes of properties, the Fregean contents of, e.g., perceptual experiences and perceptual beliefs or judgments can be modeled instead as functions that take scenarios—i.e. epistemically possible worlds “centered” on the viewpoint of a subject at a time and location—as inputs and yield extensions (properties, objects, truth values) as outputs (Chalmers 2004). Alternatively, we can think of the contents of a perceptual experience, belief, or judgment as being composed of the properties, and perhaps objects, that are presented by the experience: these entities provide a condition on accuracy, since the experience/belief/judgment will be accurate just in case the object presented in the experience has the properties it appears to have, or just in case *some* object, at a particular location relative to the perceiver, has these properties.

presentation or possible worlds. But the main argument of the chapter does not hinge on this characterization of content.

Cognitive phenomenal content must be conceptual content, since thoughts are composed of concepts, as we discussed previously. In our defense of P3 of the above argument, it will be useful for us to focus on *lexical* concepts: mental representations whose 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 the constituents of thoughts, they can be combined to form complex representations with propositional content.

We are justified in focusing on lexical concepts rather than “complete,” propositional thoughts, because the compositionality of thought means that a thought’s truth-value depends on what entities are referred to by the lexical concepts that are the thought’s constituents. Thus, while there may be *additional* challenges that confront the Standard Account of C-PI when we turn our attention to how cognitive phenomenology could be a form of propositional content—how can the phenomenology be sensitive to the difference between *John loves Mary* and *Mary loves John*, for instance—our focus on sub-propositional contents 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.¹²¹

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 is understood to include not simply the relations between the

¹²¹ For some consideration of specific problems that are raised for PI theories by propositional PI content, see Pitt (2009) and Bailey & Richards (2014).

contents of the various simpler concepts, but also a specification of these simpler contents, or “features” (Laurence and Margolis 1999, 4-5). For this reason, two complex concepts might share a *relational* structure, and yet if they are distinct concepts, then they will differ in conceptual structure by differing in the identity of some of their features.

A concept’s conceptual structure is what individuates the concept’s content, i.e. distinguishes it from all other conceptual contents.¹²² The Standard Account of C-PI says that the awareness of qualitative character of a CP experience is identical to the awareness of a thought’s conceptual content (in a cognitive-phenomenal manner), such that the individual conceptual content is part of what is *presented* in phenomenal awareness, part of what it is like for the subject. 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 subject in the qualitative character of her cognitive phenomenal experience. Presumably this must include (a) the *structural* (i.e. abstract relational) properties of that conceptual content as well as (b) those various individual constituents, or features, of the content.

A number of different theories of conceptual structure have been the subject of considerable controversy in the literature on concepts over the last several decades.¹²³ By considering each of the main candidates in turn and showing how the Standard Account for C-PI would fail given each of them, I will mount my defense of P3: SA is incompatible with

¹²² Note: to say that the conceptual structure of a concept serves to individuate its content is *not* to say that capturing the abstract structure that characterizes only how the constituents of the content relate to each other is sufficient to individuate it. Somewhat confusingly, there are two senses of ‘structure’ at use in the discussion.

¹²³ See the articles collected in Margolis and Laurence 1999 for a comprehensive introduction.

how a concept's content is individuated by its conceptual structure (and thus SA is probably false).

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 (and anyway, are unlikely to fit well with C-PI).

Arguably, C-PI requires the containment model of conceptual structure. For Awareness and Transparency require, in effect, that every factor that contributes to the CP content of a thought is "contained" within the phenomenology of the experience, in the sense that it must make a what-it's-like difference, it must show up in phenomenal awareness. Undergoing the experience is sufficient to represent the content and thus to deploy the concept. But if a concept's structure is what specifies its (narrow) semantic

properties, then the structure will need to be contained in phenomenal awareness, and every conscious tokening of the concept will include tokening its complete structure in CP experience.

To illustrate 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 reliable causal 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 capture 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.¹²⁴

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 structure of a lexical concept consists in certain relations that it stands in to other concepts, relations that 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

¹²⁴ There are tricky issues here related to phenomenal externalism (Lycan 2001; Tye 1998, forthcoming), which I leave to one side, in large part because the proponents of Strong CP and C-PI typically are internalists about phenomenology. For compelling reasons to reject phenomenal externalism, see Levine 2003a, 2003b and Pautz 2006, 2013.

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 the concept HUSBAND, had theory structure, but also that its content was fully 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 the content *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 of the concept HUSBAND need not "contain" the contents *male* and *married* as objects of phenomenal awareness in this way. Rather, this phenomenal character gets to count as the content, *husband*, just in case what one is phenomenally aware of when one is entertaining this concept, undergoing this CP experience, determines that one is disposed to draw certain inferences from this concept, i.e. those very inferences that constitute the theory structure of the concept. But, second, *how* is this phenomenal character supposed to determine that the concept stands in the requisite inferential relations to other, entirely 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 features, *male*, *married*, and so on, is *somehow* inherent in the awareness of *husband*. For it is hard to imagine what *other* aspect of this content, immediately presented in phenomenal consciousness, could be such that attending to it necessarily disposed a subject to make the requisite inferences, unless it just boiled down to “seeing” that husbands are married males. But that would be to fall back on the containment model again.¹²⁵

Of course, I am not denying, in general, 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 typical Theory theorist to explain how concepts have the conceptual 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 question why the particular CP characters for these concepts couldn’t be separated from the 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

¹²⁵ Understanding the Theory-Theory along the lines of the containment model would hardly help C-PI, given that theory structure is often thought to be holistic, in the sense that (much or all of) the theory is potentially implicated in fixing the content of any single concept. Thus, for the Standard Account of C-PI to be true, an experience for a single lexical concept would have to involve phenomenal awareness of the entire theory in which the concept is embedded!

structure, we are better off considering a different version of conceptual structure that might fit better with SA.

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

III. The Problem of Phenomenal Explosion

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).¹²⁶ For simplicity's sake, in what follows let us focus on definitional structure. According to the Definitional Theory, a concept's structure constitutes a definition of that concept, where this amounts to necessary and sufficient conditions for the correct semantic application of that concept, i.e. for something to count as falling within the extension of the concept. Since on the most natural interpretation of the theory (the containment model), 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 HUSBAND must involve these two components—e.g. categorizing something as a husband quite literally involves categorizing it as male and as married—the concept cannot be tokened without also tokening the constituent concepts MALE and MARRIAGE.

¹²⁶ For instance, Peacocke 1992, and Rosch and Mervis 1975.

Combined with C-PI, this view entails that the cognitive phenomenal experience that represents the conceptual content *husband* includes CP experiences as components that represent the conceptual contents *male* and *marriage*. Furthermore, according to the Definitional Theory, the concepts of MALE and MARRIAGE presumably are also concepts with structure, and decompose into other concepts, which constitute necessary and sufficient conditions for their application. At some point, this structural decomposition will have to terminate in primitive, unstructured concepts, but it is unlikely that such apparently complex concepts as MALE and MARRIAGE, or any of their immediate constituents, will turn out to be primitive.

According to Awareness and Transparency, conceptual contents are represented in CP experiences by forming part of the what-it's-like character of which the subject is phenomenally aware. So, on the Standard Account of C-PI, the definitional structure of these constituent concepts—and the structure of their constituents, and so on—must also be phenomenally presented by qualitative properties contained within the overall CP experience for the concept HUSBAND. 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. Call this challenge to C-PI the problem of *Phenomenal Explosion*.

The problem is not only that the entire definitional structure of a given concept would have to be *determined* by some aspect of the phenomenology of a cognitive

experience every time we entertain a thought with that content. That much is guaranteed just by C-PI, which says that a CP experience is identical to (and therefore is sufficient for) the representation of conceptual content. But once we add the Awareness and Transparency theses, the view also requires that the subject is *directly aware* of all those properties contained within the concept's definition, and the entire relational structure of those properties must be mirrored in the various discernible ways that the myriad qualitative properties are presented as relating to each other in phenomenal awareness. So the problem of Phenomenal Explosion is the challenge for C-PI of having a huge number of properties, and their relational structure, *presented* at once by the phenomenology of a single CP experience. A moment's consideration suggests that the idea of cognitive phenomenology presenting the complete definitional structure of each concept the subject entertains, each time she entertains it, is absurd.

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 very great deal of phenomenal-intentional variety and complexity, presented to our conscious awareness at every waking moment. Since this degree of qualitative richness in what it's like for the subject is not only possible but also manifestly *actual* in the case of conscious perception, there is no reason to think something similar could not be true in the case of conscious thought.

An initial reply to this point is that our experience of our thoughts does not seem anywhere near as rich and various as a typical perceptual state (does your split-second

awareness of the content *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 SA entails 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, when combined with the containment model of conceptual structure, demands that cognitive phenomenal awareness contains levels of complexity not required of perceptual awareness. To better appreciate the distinct difficulty Phenomenal Explosion poses for the Standard Account of C-PI, we can compare SA with how an analogous Phenomenal Intentionality account might work in the case of perceptual experiences.

Start with cognitive experience. Consider entertaining a thought about the nature of aristocracy. According to the SA, whatever sets an aristocracy apart from an oligarchy or plutocracy—not to mention from wildly different things like snow cones, sea turtles, and Special relativity—must be phenomenally presented in the CP experience for *aristocracy*, or else there would be nothing to make it the case that the subject was thinking about one rather than the other. If, for instance, what distinguishes aristocracy from oligarchy is that the ruling group must be a hereditary class, then there must be some qualitative property in phenomenal awareness, some aspect of what it is like for the subject thinking about aristocracy, that is constituted by *hereditariness*. Further, the relations that hold between hereditariness and e.g. holding political power, or being a socio-economic class of people—constitutive elements of the property of being an aristocracy, which relate in distinct ways to each other—must be mirrored in the relations between the phenomenal aspects of the experience that correspond to these features. Otherwise, there would be no way to account

for how attending to what the phenomenal features of experience are like could provide essential, individuating information about the content of our thought, e.g. that it has to do with a form of government in which power is held by a hereditary ruling class.

Nor are the relations between elements of the definitional structure that need to be captured in the phenomenology limited to relations that explicitly form part of the content: consider the concept, *orchestra*, the definition of which, suppose, includes the notions of *violinist*, *cellist*, and so on. The relation of playing-a-different-instrument-than, which holds between violinists and cellists, is not explicitly part of the definitional structure of the concept *orchestra*. But this relation still must be captured by 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 phenomenal character of the experience, since it is entailed by those features. By analogy: my visual experience of a group of basketball players includes the visual phenomenal presentation of their respective heights. But the phenomenology of the experience also captures the fact that, in virtue of their heights, 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 the content of a concept (i.e. the features that contribute to its definition) and the qualitative properties of that concept’s CP character, such that every relation (that makes for genuine similarity and difference) that holds essentially between the properties represented in the concept’s content maps onto either an introspectively discernible relation holding between the

qualitative properties presented in phenomenal awareness, 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 qualitative properties presented in phenomenal awareness.¹²⁷

Now, in the case of *perceptual* experience, by comparison, it seems intuitively plausible that there always will be an isomorphism between, on the one hand, the set of elements consisting of perceptual phenomenal properties that constitute what it’s like for the subject at that moment, and, on the other hand, the set of elements consisting of all those represented features of the subject’s environment that make up the perceptual content.¹²⁸ For example, if I have an experience of pain in my chest composed of three consecutive qualitative properties of which I am phenomenally aware, the content of this experience will include the positing of 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 phenomenal properties are experienced as standing in to each other. Furthermore, if the second pain quality seems more like the first than the third—e.g. because 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.

In general, if the worldly properties in the content of a perceptual experience are represented as relating to each other in some way, there will be a corresponding relation holding between qualitative properties of which the perceiving subject can become aware

¹²⁷ 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). The qualifier, “that make for genuine similarity and difference,” is intended to rule out highly unnatural relations such as being-more-like-my-grandma’s-favorite-property.

¹²⁸ It may be that this isomorphism holds in virtue of the two sets of elements being identical.

when she attends to the phenomenal character of her experience. So for complex perceptual phenomenology to be (partly) constituted by the various components of complex perceptual contents in the way prescribed by a perceptual version of the Standard Account, necessarily relational structure will be preserved between content and phenomenology.¹²⁹

However, the challenge of preserving the relational structure of the properties represented by a perceptual experience within the phenomenology of the experience is eased considerably by the fact that those properties of physical objects in our environment that we perceive—corresponding to our commonsense notions of colors, shapes, sounds, and so on—naturally form *families* of properties, 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 phenomenal properties are necessarily *homomorphic* to the families of perceptible properties of physical objects that are represented by perceptual experiences:

The distinctive qualities by means of which we classify sensations form families of properties that pertain to color, visual shape, sound, and so forth. 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. For example, the red sensory quality of visual sensations resembles the orange quality of such sensations more than either resembles the sensory green or blue of such sensations...it is the whole family of mental color properties that corresponds, by virtue of the relations that hold among its members, to the family of

¹²⁹ The role structure preservation plays in individuating the contents of perceptual experience finds a striking parallel in the way images and other *iconic* or *pictorial* representations represent their contents. For insightful discussion on this important issue, see Kulvicki 2014, and Greenberg MS, chapter 4. In a way, this chapter's whole critique of the C-PI view can be summed up in the observation that, whereas perceptual experiences (very roughly) can represent the way that pictures do, thoughts cannot. But (regardless of whether or not its proponents realize it) the Standard Account seems to require this kind of pictorial view of cognitive-phenomenal representation, so it must be false.

physical color properties... Similar observations hold for other properties of shape and other sensory modalities. (Rosenthal, 1991, 22-23)¹³⁰

The members of each family of perceptible properties 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 (phenomenally determined) contents of our perceptual experiences are ideally suited to be presented in conscious awareness as complexes of qualitative properties, since, in virtue of what they're like for the subject, sensory qualities group naturally into families, the members of which can be ordered into quality spaces.

In contrast, most of the worldly properties for which we have lexical concepts do *not* naturally or 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 David Rosenthal has noted (by "intentional content" here, he means conceptual content),

Homomorphism theory fixes qualities by reference to families of similarities and differences among properties perceptible by particular modalities. But there are no such families to use in fixing intentional content. 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

¹³⁰ A *homomorphism* is any (partial) structure-preserving map between two sets of things, of which an *isomorphism* is a more demanding type (every instance of isomorphism is also an instance of homomorphism, but not vice versa). The difference between the two does not matter for our purposes here.

concept of; so we can't appeal to such relations to fix the concept of a dog. But maybe we can appeal instead to the ways dogs resemble and differ from things within some suitable family of things, such as mammals or animals. If so, commonsense taxonomic schemes might enable homomorphism theory to apply to conceptual content after all...But even if some commonsense concepts do fit into such taxonomic schemes, many concepts do not. Concepts often apply to things whose natures fall into no such well-defined families. And the commonsense concepts that do [fall into such families], typically fit into alternative taxonomic schemes, and so would deliver conflicting sets of similarities and differences. (2006, 208)¹³¹

Let me be clear: unlike Rosenthal, I am not advocating that we attempt to *fix* the content of our concepts and thoughts purely by appeal to the various relations that the features of the content stand in to each other. My point is that the proponent of SA still needs to explain how the relational structure of a concept's content can be *matched* by what is presented to the subject in phenomenal awareness—even if she does not think that matching this relational structure would be sufficient for the qualitative properties of phenomenal awareness to count as presenting the concept's content. This passage from Rosenthal makes the point that the relational structure of most lexical concepts—unlike the relations holding between those sensible properties presented in perceptual experience—is unhelpfully complicated and messy.

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 relational structure 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

¹³¹ Rosenthal is focused here on homomorphisms only as they figure in his Higher-Order Thought theory of phenomenal consciousness, and is not concerned with cognitive phenomenology or phenomenal intentionality.

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.

The point is not that all the relations that hold between violins and trumpets, for example, are part of the content-individuating definitional structure of the corresponding concepts. The point is rather that, for a great many of these relations, their holding between these two types of instruments is entailed by the instruments' definitions (let us suppose), and so must be *retrievable* by attending to the qualitative properties presented to phenomenal awareness (the way one might notice that one basketball player stands in the taller-than relation to another player, based only on visually experiencing their respective heights). So, in contrast to the perceptual representation of colors or sounds, the whole complexity of a conceptual content's definitional structure—and the relations it stands in to other contents in virtue of that structure—in all likelihood could not be captured in phenomenal awareness by only a handful of phenomenal primitives differing along only a handful of qualitative dimensions.

Therefore, CP experiences that present the definitional structure of our lexical concepts would not *just* contain lots of distinct phenomenal properties—the way an experience of distinct fine-grained shades of blue would count as containing many distinct phenomenal properties. But part of the challenge lies in the fact that many of these CP properties would have to be very different from each other, standing in many different types of discernible relations to each other within a single experience. Thus, not only must the CP for a single concept or thought contain an explosion of phenomenal complexity in terms of the sheer number of distinct phenomenal types combined into one experience, but also in terms of the number and diversity of ways they must be 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 speaking *impossible* for cognitive phenomenology. But I am saying the idea that we have such experiences every time we consciously entertain a thought is just not credible. And given the limitations on consciousness and cognitive awareness in the human mind, I would think it's impossible that we would ever have such experiences. This is the challenge posed for C-PI by Phenomenal Explosion.

3. Someone might suggest that we can avoid this difficulty by appealing to a principle of "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 the conceptual content *aristocracy* from the content *oligarchy*.¹³² What matters is just that the political theorist's CP experiences can distinguish these two contents, and that my CP experiences simply defer somehow to the understanding of the experts.

I have two responses. First, the experts' thoughts would still face Phenomenal Explosion, which as a description of how *anyone* thinks is still hugely implausible (and a great many people have expert knowledge of *some* domains). Second, we can concede that the conceptual structure of many or most 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 view on which the individuating conceptual structure of at least a great many of our concepts is at least somewhat complex.

¹³² Thanks to Bill Child for bringing this point to my attention.

The problem could be avoided if the conceptual structure of the majority of our concepts consisted simply in representing an appeal to “whatever the experts say about x”. But if this kind of passing the buck was all that conceptual content amounted to, then it would defeat the purpose both of positing conceptual structure in the first place, and of the idea that conscious thinking amounts to experiences of “holding before one’s mind” specific thought contents (Horgan & Tienson 2002). For surely no C-PI theorist thinks that our awareness of the conceptual contents of our thoughts is exhausted by a string of introspectively indistinguishable appeals to what the experts think about x, y, and z. In that case, better just to say that concepts entirely lack structure; but we have already seen why Conceptual Atomism does not fit with the Standard Account of C-PI.

IV. The Problem of Structural Inadequacy

1. Moreover, even if we had complex CP experiences for every concept, the qualitative features of which were isomorphic to that concept’s definitional structure, this would not be enough to meet the requirements of the Standard Account. For awareness of the abstract structural properties of a content would not reveal to the introspecting subject what her thoughts were about to any interesting degree. But the Awareness thesis says that the phenomenal awareness of a CP experience is identical to the cognitive awareness of a thought’s content.

This point leads us to our second argument that the Standard Account is incompatible with a containment model of conceptual structure: for the phenomenal character of CP experiences merely to preserve the *relational* structure of a conceptual content is obviously *not* sufficient to fulfill Awareness and Transparency. An experience

whose phenomenology has the same structure as a concept's content is not, for that reason alone, presenting the content in phenomenal awareness. But arguably, it is not possible that a CP experience, using as raw materials a set of simple phenomenal primitives, could present significantly *more* than just the abstract relational structure of a concept in phenomenal awareness. Call this the problem of *Structural Inadequacy*.

We can present the problem as a dilemma: either all phenomenal character is at bottom purely *qualitative*, or it isn't. By calling properties "qualitative," I mean our grasp of their intrinsic natures—in the present case, our subjective grasp of what they are like—is conceptually independent from, and thus not exhausted by, structural and functional concepts, in a way that leads to an "explanatory gap" between our understanding of qualitative properties and the properties that form the basis of our everyday and scientific explanation (Levine 1982, Chalmers 1996). Throughout the thesis, I have been operating under the assumption that all phenomenal character is fundamentally qualitative in this sense, and I have been using 'qualitative' and 'qualities' as basically synonymous with the "what-it's-like-ness" for the subject. But it is a debatable question as to whether all phenomenology—all properties that contribute to what it is like to undergo an experience in the distinctive first-person sense—really are qualitative, if that means, e.g., they can be appealed to in the standard anti-materialist arguments, such as the Knowledge Argument. Perhaps some experiences just amount to being aware of ordinary, worldly properties and things, which can be fully described in commonsense and/or scientific terms without any mysterious, irreducibly qualitative remainder. Either way, I think the Standard Account will have an insuperable problem accounting for how the basic, non-relational elements

that make up a concept's content (in this case, the primitive features of its definitional structure) can be presented in phenomenal awareness.

First, suppose phenomenology is at bottom purely qualitative. Then the phenomenal primitives out of which CP experiences must be constructed are simple qualitative properties similar to colors or smells, and so cannot be suitably informative, cannot point us towards non-qualitative, everyday, physical-functional content in isolation, but only insofar as they are presented in complex configurations and stand in discernible relations to one another. For by stipulation, the intrinsic, subjective natures of these qualities do not contribute to conveying information about the physical world—except insofar as they can be discerned to be more or less similar to other qualities. But the burden on pure relational structure is too much for it to bear. Even if these relations succeed in mirroring the relations that hold between the conceptual primitives that combine to form a concept's content, they still will fail in many cases to present that content to the subject, because the conceptual primitives themselves—the contents of the unstructured atomic concepts out of which all other concepts are built—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, ineffable quality cannot capture.¹³³

In the case of *perceptual* experience, we can see how qualitative primitives can be combined to present useful information to the perceiver about her environment. Recall Rosenthal's earlier contention that our commonsense understanding of perceptible

¹³³ We could propose that every conceptual primitive is presented in phenomenal awareness by a complex of simple qualities (or perhaps an irreducibly complex quality?), thus introducing more phenomenal structure to match the internal structure of the conceptual primitives. But where would it end? The world according to contemporary physics might turn out to be structure all the way down to the ground floor, but surely the structure of our ordinary concepts does not reach down that far. CP character does not make us phenomenally aware of all the world's structure, right down to the building blocks of reality at the quantum level, every time we entertain a normal, macro-level concept!

properties such as the physical colors essentially makes reference to their similarities and differences. This commonsense understanding is explained by the fact that the systematic relations of similarity and difference that hold between (families of) perceptible properties form part of the narrow, phenomenal content of sensory-perceptual experiences, content that is then taken up into perceptual beliefs and used to pick out the perceptible properties in the subject's environment. For the qualitative character of perceptual experience presents properties to the conscious subject that are (partly) individuated by those very relations of similarity and difference in which the perceptible properties stand to each other. Awareness of these properties determines a kind of narrow content. And this narrow content, when combined with facts about the subject's physical environment, then yields the experience's wide content.¹³⁴ So it's plausible to suppose that our grasp of the perceptible properties comes partly via our grasp of the systematic ways they relate to each other.

By contrast, the complicated relational structure of a lexical concept's content, *when separated from facts about the intrinsic natures of the fundamental elements of this content*, does not typically play a role in our conscious grasp of it. Certainly the proponents of C-PI do not have a set of purely relational properties in mind as thought contents when they speak of phenomenology determining the contents of thought. Rather, the phenomenal contents of thought are supposed to consist of the sorts of properties that internal duplicates Oscar and Twin Oscar can both think about, in spite of differences in their respective physical environments: both Oscar and his twins have "water-y" concepts, "rabbit-y" concepts, "gold-y" concepts, and so on. But undergoing an experience whose

¹³⁴ Cf. Chalmers 2004, 2006; Horgan 2014; Horgan & Tienson 2002.

phenomenal features are merely isomorphic to the elements of the definitional structure of Oscar's water concept is clearly *not* sufficient to have a "water-y" concept in any intuitive sense. Isomorphism is cheap: those same phenomenal features could match all sorts of *other* definitions that just so happen to share an abstract relational structure with the definition of water. So if all CP is at bottom composed of qualitative simples that can only present a concept's content to us by mirroring the content's relational structure, then the Standard Account faces the problem that presenting relational structure is not sufficient for a CP experience to present the rich totality of conceptual content.

2. Suppose instead that certain fundamental CP properties are *not* qualitative in the above sense, but are identical to certain real-world physical, functional, or structural properties, and are presented in experience as such. We might think that the problem of Structural Inadequacy dissolves. For in that case, according to Transparency, what it's like to undergo an experience with that CP property just consists in being phenomenally aware of it. If so, then although cognitive experiences are not able to present a concept's content simply by mirroring its structure, that does not matter, since individual primitive phenomenal properties can present the various constituents of that content.

Being phenomenally aware of a conceptual content still will involve being aware of all the properties and relations that contribute to its complex conceptual structure. But the conceptual content of a CP experience will not have to be derived by constructing it from the similarity and difference relations holding between purely qualitative building blocks, since awareness of conceptual features is intrinsic to the phenomenal ground floor, so to speak.

However, even if the phenomenal primitives that combine to form a CP experience for a lexical concept are not pure, simple qualities, but rather properties found in the external world, these properties still have to be utterly simple, with little or no internal variation or complexity. For any complexity that plays a role in individuating the content must be transparently presented in phenomenal awareness. So on pain of infinite regress, the properties and relations in phenomenal awareness must ultimately decompose into simple phenomenal atoms and simple relations. And that leads again back to *structure*—to ways that the simple, largely uninformative phenomenal primitives combine and relate—as the only means by which to phenomenally determine the great variety of our complex conceptual contents.

Furthermore, it is very hard to think of any worldly properties that plausibly could play the role of these phenomenal primitives. In any case it seems highly unlikely that all of our lexical concepts could be constructed by somehow relating such utterly simple conceptual primitives to each other. The best candidates for the role seem to be the sensible properties represented by sensory-perceptual experiences (perhaps along with a few others, such as objecthood and causation). But the empiricist tradition, which sought to base all human concepts on our ideas of sensible properties, does not have a particularly happy history. And anyway, to make the phenomenal atoms of conceptual experiences sensory phenomenal properties would be to abandon the need for a *sui generis cognitive* type of phenomenology.

In summary, the problem of Structural Inadequacy is a challenge for SA regardless of whether cognitive phenomenology is at bottom qualitative or not. This argument, in conjunction with the argument from Phenomenal Explosion, establishes that if our lexical

concepts have conceptual structure according to the containment model, then CP experiences cannot fit the description of them offered by the Standard Account of C-PI.

3. It should be clear that other accounts of conceptual structure that are 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). Prototype structure swaps the necessary and sufficient conditions of definitions for representations of properties that are not *necessarily* but only *probably* instantiated in members of the category picked out by the concept, and which do not need to be *exhaustively* but only *largely* instantiated in order to be sufficient to determine that something belongs to the category. But if prototype structure consists in complete sets of prototypical features for each concept, which are contained in the concept itself, then this theory cannot avoid the problems that confront the Definitional view. For all the features that are elements of a concept's prototype structure will need to be present in every token CP experience for that concept, or else the phenomenology of the experience will not be sufficient to capture the phenomenal content of the concept and, together with wide environmental facts, individuate the concept as the type that it is.

It is true that individuals may belong in the extension of a concept without instantiating all the properties included in the concept's prototype structure, and thus that a subject may recognize that an individual falls under the concept without including all the concept's prototypical features in its representation of the individual. But the prototype structure, including all its features and their statistical weightings, somehow must be part

of the C-PI experiences that represent the concept's content, which means all these features—and the features that constitute their structure, and so on—must be presented in phenomenal awareness. And then the problems of Phenomenal Explosion and Structural Inadequacy apply in full force, once again. Yet if we reject the containment model of prototype structure altogether in favor of an inferential model, then, as I argued above, it seems a concept's content would not be exhausted by what is found in phenomenal awareness alone. The inevitable conclusion to draw is that CP experiences are inadequate vehicles for capturing a lexical concept's complete conceptual structure, and thus inadequate representations of (even a kind of narrow) conceptual content.

V. Concluding Remarks on the Chapter and the Thesis

This completes my defense of Premise 3 of my argument against the Standard Account of C-PI, my argument for why, in all likelihood, cognitive phenomenology cannot determine the content of our thoughts. Recall that P3 was the following claim:

P3 But the Standard Account—C-PI, Awareness, and Transparency—is incompatible with all known viable theories of conceptual structure

From this claim, I conclude that the Standard Account is probably false and should be rejected, and from there, that C-PI and the Modal Thesis should be rejected too. So no phenomenal property is sufficient to metaphysically determine what, or that, I am thinking.

We might want to resist this line of thought, on the grounds that, for all we know, the *true* theory of conceptual structure will 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 arguments 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 atomistic alternative. Indeed, it is difficult to see how any concept *could* have internal structure that was not connected 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 involves a subject's direct conscious awareness of a thought content, "holding it before one's mind" (Horgan and Tienson 2002). 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 *way too much structure* for us to capture it in normal human phenomenal awareness. So the prospects of combining C-PI with a realistic view of thought contents looks dim.

Where does all this leave the cognitive phenomenology debate? I have suggested that without C-PI, or something like it, it is difficult to see why anyone would accept the Modal Thesis, and I have just argued that C-PI is probably false. Earlier, we saw that there are no clear a priori or a posteriori reasons for supposing that thoughts cannot be experiences in the stream of consciousness, i.e. instantiate their own phenomenal properties. And we saw that standard contrast case arguments in the literature provide some evidence for believing that experiences of conscious thinking involve phenomenal character over and above that of sensory experiences. The question then becomes, what type of CP best explains the evidence? If I am right that the Modal Thesis without C-PI is a

nonstarter, then the majority of Phenomenal Liberals, who have embraced Strong CP as best capturing what it's like to think, have been appealing to the wrong explanation. We are better off exploring the prospects of a robust Moderate CP theory, a view which allows that there may be far more to the phenomenal mind than low-level sensory-perceptual experiences and imagery, but which denies that what it's like to think is ever really like holding a thought content suspended before the mind's eye. I have proposed that William James's idea of fringe phenomenology is a promising starting point from which to begin developing this robust Moderate CP account.

Bibliography

Anscombe, G.E.M., and P.T. Geach. *Three Philosophers*. Oxford University Press, 1961.

Aydede, Murat. "The Language of Thought Hypothesis." In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Fall 2010.
<<http://plato.stanford.edu/archives/fall2010/entries/language-thought/>>.

Bailey, Andrew R., and Richards, Bradley. "Horgan and Tienson on Phenomenology and Intentionality." *Philosophical Studies* 167, no. 2 (2014). 313-326.

Barsalou, Lawrence. "Perceptual Symbol Systems." *Behavioral & Brain Sciences* 22 (1999). 577-660.

Bayne, Tim. "Perception and the Reach of Phenomenal Content." *Philosophical Quarterly* 59, no. 236 (2009). 385-404.

—. *Thought: A Very Short Introduction*. Oxford University Press, 2013.

—. "The Puzzle of Cognitive Phenomenology", MS.

Bayne, Tim, and David Chalmers. "What Is the Unity of Consciousness?" In *The Unity of Consciousness: Binding, Integration, and Dissociation*, edited by Axel Cleermans. Oxford University Press, 2003.

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

Block, Ned. "Inverted Earth." *Philosophical Perspectives* 4 (1990). 53-79.

—. "On a Confusion about a Function of Consciousness." *Behavioral and Brain Sciences* 18 (1995). 227-287.

—. "Seeing-As in the Light of Vision." *Philosophy and Phenomenological Research* 89, no. 3 (2014). 1-13.

Bor, D., J. Duncan, and A.M. Owen "The Role of Spatial Configuration in Tests of Working Memory Explored with Functional Neuroimaging." *Scandinavian Journal of Psychology* 42, no. 3 (2001). 217–224.

Bor, Daniel, John Duncan, Andy C.H. Lee, Alice Parr, and Adrian M. Owen. "Frontal Lobe Involvement in Spatial Span: Converging studies of normal and impaired function." *Neuropsychologia* 44, no. 2 (2006). 229–237.

Briscoe, Robert. "Mental Imagery and the Varieties of Amodal Perception." *Pacific Philosophical Quarterly* 92 no. 2 (2011). 153–173.

—. "Vision, Action, and Make-Perceive." *Mind and Language* 23, no. 4 (2008). 457–497.

Buckner, R. L., W.M Kelley, and S.E. Petersen. "Frontal Cortex Contributes to Human Memory Formation." *Nature Neuroscience* 2 no. 4 (1999). 311–314.

Burge, Tyler. *Origins of Objectivity*. Oxford University Press, 2010.

Byrne, Alex. "Introspection." *Philosophical Topics* 33 (2005). 79–104.

—. "Knowing That I Am Thinking." In *Self-Knowledge*, edited by A. Hatzimoysis. Oxford University Press, 2008.

—. "Review of Kriegel, ed., Phenomenal Intentionality." *Notre Dame Philosophical Reviews*. 9 19, 2013. <https://ndpr.nd.edu/news/42908-phenomenal-intentionality/>.

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

—. "Knowledge Acquisition: Enrichment or conceptual change?" In *Concepts: Core Readings*, edited by Eric Margolis & Stephen Laurence, 459–487. MIT Press, 1999.

Carruthers, Peter. *Phenomenal Consciousness: A Naturalistic Theory*. Cambridge University Press, 2000.

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

Chalmers, David. "Perception and the Fall from Eden." In *Perceptual Experience*, edited by Tamar S. Gendler & John Hawthorne, 49–125. Oxford University Press, 2006.

—. *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 University Press, 2004.

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

—. "Intellectual Gestalts." In *Phenomenal Intentionality*, edited by Uriah Kriegel, 175-191. Oxford University Press, 2013.

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

—. "The Intentional Structure of Consciousness." In *Consciousness: New Philosophical Perspectives*, edited by Quentin Smith and Aleksandar Jokic, 33-56. Oxford University Press, 2003.

—. "Unconscious Belief and Conscious Thought." In *Phenomenal Intentionality*, edited by Uriah Kriegel. Oxford University Press, 2013.

Cummins, Robert. "'How Does It Work' versus 'What Are the Laws?': Two conceptions of psychological explanation." In *Explanation and Cognition*, by Frank Keil and Robert A. Wilson, 117-145. MIT Press, 2000.

Dolan, R.J., and P.C. Fletcher. "Dissociating Prefrontal and Hippocampal Function in Episodic Memory Encoding." *Nature* 388 (1997). 582-585.

Epstein, Russell. "Consciousness, Art, and the Brain: Lessons from Marcel Proust." *Consciousness and Cognition* 13, no. 2 (2004). 1-31.

—. "The Neuro-Cognitive Basis of the Jamesian Stream of Thought." *Consciousness and Cognition* 9 (2000). 550-575.

Evans, Gareth. *The Varieties of Reference*. Oxford University Press, 1982.

Farkas, Katalin. "Constructing a World for the Senses." In *Phenomenal Intentionality: New Essays*, edited by Uriah Kriegel. Oxford University Press, 2013.

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

—. "Burge on Perception." MS.
http://rucss.rutgers.edu/faculty/pylyshyn/Burge_OriginsOfObjectivity/Fodor_BurgeOnPerception.docx (accessed 2015).

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

—. "Connectionism and the Problem of Systematicity (continued): Why Smolensky's solution still doesn't work." *Cognition* 62, no. 1 (1997). 109-119.

—. *In Critical Condition: Polemical Essays on Cognitive Science and the Philosophy of Mind*. MIT Press, 2000.

- . *The Modularity of Mind*. MIT Press, 1983.
- . "The Revenge of the Given." In *Contemporary Debates in Philosophy of Mind*, edited by Brian P. McLaughlin and Jonathan D. Cohen, 105-116. Blackwell, 2007.
- . *The Language of Thought*. Harvard University Press, 1975.
- . "Why There Still Has to be a Language of Thought." In *Psychosemantics*. MIT Press, 1987.
- Fodor, Jerry, and Zenon Pylyshyn. "Connectionism and Cognitive Architecture: A critical analysis." *Cognition*, no. 28 (1988). 3-71.
- Frey, Christopher. "Phenomenal Presence." In *Phenomenal Intentionality: New Essays*. Oxford University Press, 2013.
- Galín, David. "Aesthetic Experience: Marcel Proust and the neo-Jamesian structure of awareness." *Consciousness and Cognition* 13, no. 2 (2004). 241-253.
- . "The Structure of Awareness: Contemporary applications of William James' forgotten concept of 'The Fringe'." *Journal of Mind and Behavior* 15, no. 4 (1994). 375-402.
- Geach, P.T. *God and the Soul*. Routledge and Kegan Paul, 1969.
- . *Mental Acts*. Routledge and Kegan Paul, 1957.
- Gray, J. A., and McNaughton, N. *The Neuropsychology of Anxiety: An enquiry into the function of the septo-hippocampal system*. Oxford University Press, 2000.
- . "The Contents of Consciousness—A neuropsychological conjecture." *Behavioral and Brain Sciences* 18, no. 4 (1995). 659-676.
- Greenberg, Gabriel. *The Semiotic Spectrum*. Ph.D. dissertation, Rutgers University, 2011.
- Harman, Gilbert. "The Intrinsic Quality of Experience." *Philosophical Perspectives: Action Theory and Philosophy of Mind* 4 (1990). 31-52.
- Horgan, Terence. "Phenomenal Intentionality and Secondary Qualities: The Quixotic case of color." In *Does Perception Have Content?*, edited by Berit Brogaard, 330-49. Oxford University Press, 2014.
- . "From Agentive Phenomenology to Cognitive Phenomenology: A guide for the perplexed." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 57-78. Oxford University Press, 2011.

Horgan, Terence E., John L. Graham, and George Tienson. "The Phenomenology of First-Person Agency." In *Physicalism and Mental Causation*, by Sven Walter & Heinz-Dieter Heckmann (eds.), 323-40. Imprint Academic, 2003.

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 University Press, 2002.

Jackendoff, Ray S. *Consciousness and the Computational Mind*. MIT Press, 1987.

James, William. *The Principles of Psychology*. Holt, 1890.

Janowsky, J. et al. "Memory and Metamemory: Comparisons between patients with frontal lobe lesions and amnesiac patients." *Psychobiology* 17 (1989). 2-11.

Kind, Amy. "What's So Transparent about Transparency?" *Philosophical Studies*, 115, no. 3 (2003). 225-244.

Koriat, Asher. "Metacognition and Consciousness." In *Cambridge Handbook of Consciousness*, edited by Philip David Zelazo, Morris Moscovitch, and Evan Thompson, 289-235. Cambridge University Press, 2007.

—. "The Feeling of Knowing: Some metatheoretical implications for consciousness and control." *Consciousness and Cognition* 9, no. 2 (2000). 149-171.

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

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

Kriegel, Uriah, and Kenneth W. Williford. *Self-Representational Approaches to Consciousness*. MIT Press, 2006.

Kulvicki, John V. *Images*. Routledge, 2014.

Laurence, Stephen, and Eric Margolis. "Concepts and Cognitive Science." In *Concepts: Core Readings*, edited by Eric Margolis and Stephen Laurence, 3-81. 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 (1983). 354-61.

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

Loar, Brian. "Phenomenal Intentionality as the Basis of Mental Content." In *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, edited by M. Hahn and B. Ramberg. MIT Press, 2003.

Lormand, Eric. "Nonphenomenal Consciousness." *Nous* XXX (1996). 242-261.

Lycan, William. "The Case for Phenomenal Externalism." *Philosophical Perspectives* 15, no. s15 (2001). 17-35.

Mangan, Bruce. *Meaning and the Structure of Consciousness: An Essay in Psycho-Aesthetics*. 1991.

—. "Sensation's Ghost: The non-sensory "fringe" of consciousness." *Psyche* 7, no. 18 (2001).

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

Marr, David. *Vision*. Freeman, 1982.

Masrour, Farid. "Is Perceptual Phenomenology Thin?" *Philosophy and Phenomenological Research* 83, no. 2 (2011). 366-397.

—. "Phenomenal Objectivity and Phenomenal Intentionality." In *Phenomenal Intentionality*, edited by Uriah Kriegel. Oxford University Press, 2013.

Matthen, Mohan. "Two Visual Systems and the Feeling of Presence." In *Perception, Action, and Consciousness: Sensorimotor Dynamics and Two Visual Systems.*, edited by Nivedita Gangopadhyay, Michael Madary and Finn Spicer. Oxford University Press, 2010.

McGovern, Katherine. "Feelings in the Fringe." *Consciousness and Cognition* 2, no. 2 (1993). 119-125.

Montague, Michelle. "The Life of the Mind." In *Phenomenal Qualities: Sense, Perception, and Consciousness*, edited by Paul Coates and Sam Coleman. Oxford University Press, forthcoming.

Moore, G.E. "The Refutation of Idealism." *Mind* 12, no. 48 (1903). 433-453.

Moscovitch, Morris, and Gordon Winocur. "The Frontal Cortex and Working with Memory." In *Principles of Frontal Lobe Function*, 188-209. 2002.

Nanay, Bence. "Action-Oriented Perception." *European Journal of Philosophy* 20, no. 3 (2012). 430-446.

—. "Do We See Apples as Edible?" *Pacific Philosophical Quarterly* 92, no. 3 (2011). 305-322.

Nes, Anders. "Thematic Unity in the Phenomenology of Thinking." *The Philosophical Quarterly* 62, no. 246 (2012). 84-105.

O'Keefe, J. and Nadel, L. *The Hippocampus as a Cognitive Map*. Oxford University Press, 1978.

O'Shaughnessy, Brian. *Consciousness and the World*. Oxford University Press, 2000.

—. "The Epistemology of Physical Action." In *Agency and Self-Awareness*, edited by Johannes Roessler and Naomi Eilan. Clarendon Press, 2003.

Pautz, Adam. "Does Phenomenology Ground Mental Content?" In *Phenomenal Intentionality: New Essays*, edited by Uriah Kriegel, 194-234. Oxford University Press, 2013.

—. "Sensory Awareness Is Not a Wide Physical Relation: An empirical argument against externalist intentionalism." *Nous* 40, no. 2 (2006). 205-240.

—. "What Are the Contents of Experiences?" *The Philosophical Quarterly* 59, no 236 (2009). 483-507.

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

—. "Precis of a Study of Concepts." *Philosophy and Phenomenological Research* 56 (1996). 407-411.

Pitt, David. "Intentional Psychologism." *Philosophical Studies* 146, no. 1 (2009). 117 - 138.

—. "Introspection, Phenomenality, and Content." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 141-173. Oxford University Press, 2011.

—. "Mental Representation." In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Fall 2013.

<http://plato.stanford.edu/archives/fall2013/entries/mental-representation/>.

—. "The Phenomenology of Cognition Or What Is It Like to Think That P?" *Philosophy and Phenomenological Research* LXIX (2004). 1-36.

Price, Mark C. "Measuring the Fringes of Experience." *Psyche* 8, no. 16 (2002).

Price, Mark, and Elisabeth Norman. "Intuitive Decisions on the Fringes of Consciousness: Are they conscious and does it matter?" *Judgement and Decision Making* 3, no. 1 (2008). 28-41.

- Prinz, Jesse. *Furnishing the Mind: Concepts and their Perceptual Basis*. MIT Press, 2002.
- . *Gut Reactions: A Perceptual Theory of the Emotions*. Oxford University Press, 2004.
- . "The Sensory Basis of Cognitive Phenomenology." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 174-196. Oxford University Press, 2011.
- Pylyshyn, Zenon W. *Seeing and Visualizing: It's Not What You Think*. A Bradford Book, 2003.
- Rashbrook, Oliver. "Individuating Experiences and the Phenomenal Unity of Consciousness." unpublished manuscript.
- Reber, Rolf, Pascal Wurtz, and Thomas D. Zimmermann. "Exploring "Fringe" Consciousness: The subjective experience of perceptual fluency and its objective bases." *Consciousness and Cognition* 13, no. 1 (2004). 47-60.
- Robinson, William S. "A Frugal View of Cognitive Phenomenology." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague. Oxford University Press, 2011.
- . "Thoughts without Distinctive Non-imagistic Phenomenology." *Philosophy and Phenomenological Research* LXX (2005). 534-561.
- Rosch, Eleanor, and Carolyn Mervis. "Family Resemblances: Studies in the internal structures of categories." *Cognitive Psychology* 7 (1975). 573-605.
- Rosenthal, David M. "Sensory Qualities, Consciousness, and Perception." In *Consciousness and Mind*, by David M. Rosenthal, 175-226. Oxford University Press, 2006.
- . "The Independence of Consciousness and Sensory Quality." *Philosophical Issues* 1 (1991). 15-36.
- Schnyer, David M., Lindsay Nicholls, and Mieke Verfaellie. "The Role of VMPC in Metamemorial Judgments of Content Retrievability." *Journal of Cognitive Neuroscience* 17, no. 5 (2005). 832-846.
- Shimamura, A. P. "Memory and Frontal Lobe Function." In *The Cognitive Neurosciences*, by M. S. Gazzaniga (ed.). MIT Press, 1995a.
- . "Memory and the Prefrontal Cortex." In *Structure and Functions of the Human Prefrontal Cortex*, edited by K. J. Holyoak, F. Boller, and J. Grafman. New York Academy of Sciences, 1995b.
- Siegel, Susanna. "Subject and Object in the Contents of Visual Experience." *Philosophical Review* 115, no. 3 (2006b). 355-388.

- . "The Contents of Perception." *The Stanford Encyclopedia of Philosophy*. Edited by Edward N. Zalta. 2015.
<http://plato.stanford.edu/archives/spr2015/entries/perception-contents>.
- . "Which Properties Are Represented in Perception?" Vol. 84, in *Perceptual Experience*, edited by Tamar S. Gendler and John Hawthorne, 481-503. Oxford University Press, 2006a.
- Siewert, Charles. "Is Experience Transparent?" *Philosophical Studies* 117 (2004). 15-41.
- . "Phenomenal Thought." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 236-267. Oxford University Press, 2011.
- . *The Significance of Consciousness*. Princeton University Press, 1998.
- Smithies, Declan. "The Nature of Cognitive Phenomenology." *Philosophy Compass* 8, no. 8 (2013a). 744-754.
- . "The Significance of Cognitive Phenomenology." *Philosophy Compass* 8, no. 8 (2013b). 731-743.
- Soteriou, Matthew. "Content and the Stream of Consciousness." *Philosophical Perspectives* 21 (2007). 543-568.
- Spener, Maja. "Mind Independence and Visual Phenomenology." In *Introspection and Consciousness*, edited by Declan Smithies and Daniel Stoljar. Oxford University Press, 2012.
- Squire, L. R., and S. Zola-Morgan. "The Medial Temporal Lobe Memory System." *Science* 253 no. 5026 (1991). 1380-1386.
- Stich, Stephen. "Beliefs and Subdoxastic States." *Philosophy of Science* 45 (1978). 499-518.
- Stokes, Dustin. "Cognitive Penetrability of Perception." *Philosophy Compass* 8, no. 7 (2013). 646-663.
- Stoljar, Daniel. "The Argument from Diaphanousness." Edited by M. Escudria, Robert J. Stainton, and Christopher D. Viger, 341-390. *Canadian Journal of Philosophy*. University of Alberta Press, 2004.
- Strawson, Galen. "Cognitive Phenomenology: Real life." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 285-325. Oxford University Press, 2011.
- . *Mental Reality*. MIT Press, 1994.
- Topolinski, Sascha, and Fritz Strack. "Scanning the "Fringe" of Consciousness: What is felt and what is not felt in intuitions about semantic coherence." *Consciousness and Cognition* 18, no. 3 (2009). 608-618.

Tucker, Chris, ed. *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. Oxford University Press USA, 2013.

Tye, Michael. "Inverted Earth, Swampman, and Representationalism." *Philosophical Perspectives* 12, no. S12 (1998). 459-78.

—. "Representationalism and the Transparency of Experience." *Nous* 36, no. 1 (2002). 137-151.

Tye, Michael, and Briggs Wright. "Is There a Phenomenology of Thought?" In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague, 326-344. Oxford University Press, 2011.

Velleman, J. David, and Thomas Hofweber. "How to Endure." *The Philosophical Quarterly* 61, no. 242 (2011). 37-57.

Wittgenstein, Ludwig. *Philosophical Investigations*. 2nd Edition, reissued. Translated by G.E.M. Anscombe. Blackwell, 1968.