

Betz-Richman, Noah
2018 Philosophy Thesis

Title: Phenomenal Intentionality, Private Concepts, and the Problem of
Externalism:

Advisor: Joe Cruz

Advisor is Co-author: None of the above

Second Advisor:

Released Beyond Williams: release now

Contains Copyrighted Material: No

Phenomenal Intentionality, Private Concepts, and the Problem of Externalism

by

Noah Betz-Richman

Joe Cruz, Advisor

**A thesis submitted in partial fulfillment
of the requirements for the
Degree of Bachelor of Arts with Honors
in Philosophy**

WILLIAMS COLLEGE

Williamstown, Massachusetts

April 23rd, 2018

Acknowledgements

I am incredibly grateful to the Williams College Philosophy faculty for their support throughout my college career and the process of writing this thesis. I would like to thank Bojana Mladenovic for demonstrating how fun philosophy can be during my freshman fall. Since my first philosophy course, her immense support of my work and aspirations has been indispensable; she also helped shape my thoughts in the early stages of this project. Many thanks to Melissa Barry, who has challenged me to construct arguments with a high degree of rigor, and Keith McPartland, who helped focus my philosophical interests on the mind and language and first brought externalism to my attention. Finally, I would like to thank my advisor, Joe Cruz, who has supported this project since it was nothing more than an inkling in conversation behind Thompson Chapel, and who, through months of meetings, has challenged me to continually refine and sharpen my thinking. This has been one of the most rewarding intellectual experiences of my life, and I owe my deepest gratitude to all those who helped make it possible.

Table of Contents

<u>Chapter One: The Landscape of Views on Consciousness and Intentionality</u>	3
I. The Target Phenomena	3
II. The Relationship Between Phenomenal Consciousness and Intentionality	10
<u>Chapter Two: Cognitive Phenomenal Holism</u>	20
I. Introduction	20
II. Cognitive Phenomenology	24
III. Cognitive Phenomenal Holism	28
IV. Synchronic Cognitive Phenomenology	32
V. Holism	36
VI. The Dispositional Network of Beliefs	38
<u>Chapter Three: The Externalist Threat</u>	43
I. Externalism and its Motivations	43
II. PIT Responses to Externalism	50
III. Pitt's Response	56
IV. Private Concepts	58
V. Similarity of Cognitive Phenomenal Content	66
VI. Massive Theory Change	70
VII. Public Concepts	74

Chapter One: The Landscape of Views on Consciousness and Intentionality

I. The Target Phenomena

In this thesis, I outline and defend a form of Phenomenal Intentionality Theory (PIT). PIT is an approach to intentionality which aims to account for the property of intentionality in terms of phenomenal consciousness. Before beginning to discuss phenomenal intentionality theory, I will define intentionality and phenomenal consciousness. This is not a straightforward project, as ongoing debates in the philosophy of mind about the nature of phenomenal consciousness and intentionality are in some ways debates about their definitions. I will begin by trying to adhere to theory-neutral definitions, and supplement them by providing a brief landscape of views about each property.

a. Intentionality

Intentionality is the capacity of the mind to represent states of affairs. It is most straightforwardly exhibited by what are often called propositional attitudes: beliefs, desires, hopes, etc. In such states, an individual takes a certain attitude (e.g. believing) towards a certain propositional content (e.g. that snow is white). As a result, such states instantiate the representational properties of having certain intentional content. In addition to representation, intentional states are also often characterized by ‘aboutness’ or ‘directedness’ toward some content. This limited definition is meant to be neutral among views on intentionality. By examining competing approaches to understanding this property, I hope to expose a variety of intuitions about intentionality.

A dominant class of views on intentionality are causal-teleological, explaining how mental states represent states of affairs in virtue of a causal story of the relationship between that

mental state and the state of affairs it represents. These theories are naturalistic, as they aim to explain intentionality in entirely non-intentional terms. Fred Dretske and Ruth Millikan each explain representation as an evolutionarily selected biological proper function.¹ For example, an organism might possess a system which was evolutionarily selected to detect the presence of snakes. One can imagine why this system would be evolutionarily adaptive: if snakes are a natural predator, the organism will be more evolutionarily fit if it can detect snakes so as to flee from them. In Dretske's terminology, the system would act as a natural indicator of snakes, since snakes cause the system to indicate the presence of snakes.

Any theory of intentionality must to explain how a system tracks states of affairs in the world, but the system also must exhibit some independence from the states of affairs it tracks. A system like a thermometer, which is nomically connected to changes in temperature, lacks sufficient independence to exhibit intentionality (it is merely an indicator). As a result, any intentional system must possess the possibility to misrepresent the world. So long as the function of the system is to indicate snakes, cases of misrepresentation can still be explained causally in terms of the function of the system which generated the representation (the system simply failed to function properly in such cases). What is essential to such accounts is that the representation is the result of a causal relationship between a system with a representational function and a state of affairs in the external world. One implication of this approach is that these accounts are reductive: in order to understand how a system instantiates intentionality, one can investigate how that system's representational functions operate and how they were selected over time.

¹ Fred Dretske, *Knowledge and the Flow of Information* (Cambridge: MIT Press, 1981); Ruth Millikan, "Biosemantics," *Journal of Philosophy* 86 (2009): 281-297.

It is worth noting that these causal-teleological approaches to intentionality, though they may not state it explicitly, approach intentionality from a third-personal or behavioral perspective. Consider again the organism which possesses an evolutionarily adaptive system to indicate the presence of snakes. The organism's representational capacity is directly tied to a behavioral output: it represents snakes so as to flee from them. An external observer can recognize that the organism is engaging in representation only if it behaves accordingly. This behavioral approach fits neatly with the evolutionary focus of the causal-teleological theories: representation is one element in a complex set of evolutionary adaptations allowing organisms to successfully navigate their environments. Intentional states are an incredibly useful way to interpret behavior; often, we use intentional idioms to explain the behavior of systems which do not represent anything at all.

A variety of theories fall under this causal-teleological category of approaches to intentionality.² These accounts vary in the particular tracking relation they select between representations and states of affairs. However, they are united in that they see representations as bearing an essential relation to the external world in which the states of affairs represented actually occur. This intuition is strongly tied to externalism about mental content. Externalism is the view that factors in the external world play an essential role in determining mental content. As a result, intrinsically identical individuals could have relevantly divergent mental content if the circumstances in their respective environments were similarly relevantly divergent. Influential arguments by Hilary Putnam and Tyler Burge have popularized these externalist intuitions among philosophers. The distinction between wide and narrow content is useful in

² See Dennis Stampe, "Toward a Causal Theory of Linguistic Representation" (1977); Jerry Fodor, "Semantics, Wisconsin Style" (1984).

understanding the externalist intuitions. Using Chalmers' definition, "a property is *narrow* when necessarily, for any individual who has the property, an intrinsic duplicate of that individual has the property (regardless of environment)" and "a property is *wide* when it is possible for an individual to have the property while an intrinsic duplicate lacks that property."³ Externalists assert that mental content is wide, since it depends on factors external to intrinsic duplicates. The naturalist-externalist view of intentionality is that intentionality is a phenomenon which is instantiated when the right type of tracking relation between an organism and its external environment occurs, and as a result, the intentional content depends in part on the external environment.

Despite the appeal of these intuitions, there are some reasons to think that intentionality is irreducible, and that reductive naturalizing accounts may be missing something essential about the phenomenon. In particular, John Searle's Chinese Room Argument is a useful starting point to pump intuitions about the essential internality of intentionality.⁴ Searle's Chinese Room is primarily meant as an argument against the possibility of Strong Artificial Intelligence, the view that a system with the proper functional configuration of inputs and outputs would have a mind in exactly the same sense that humans have minds. However, his argument can be extended against reductive accounts of intentionality more generally. In his famous thought experiment, he imagines that he is inside a room with a rulebook that allows him to perform formal operations on Chinese symbols to produce new outputs of symbols. Crucially, he does not understand Chinese. As a result, though the room functionally behaves as if it understands Chinese, in fact

³ David Chalmers, "The Representational Character of Experience," in *The Future For Philosophy*, ed. Brian Leiter (Oxford: Oxford University Press, 2004).

⁴ John R. Searle, "Minds, Brains, and Programs," *The Behavioral and Brain Sciences* 3, no. 3 (1980): 418.

no part of the system understands Chinese.⁵ Though the system can functionally reproduce syntax, it lacks any semantics. In contrast, the man in the box can fluently understand an English input with both syntax and semantics. The functional system lacks some necessary conditions for intentionality.

Searle argues that one of the essential features of an intentional state is that it represents content under an aspect, or that it has an aspectual shape. For example, one can determinately think that the Eiffel Tower is in Paris, and not be thinking that an iron structure built in 1889 is in the capital of France, despite the fact that the terms are co-extensive. This feature is essential to intentional content: “To be intentional, a state or process must be thinkable or experienceable; and to be thinkable or experienceable, it must have an aspectual shape under which it is at least in principle, consciously thinkable or experienceable.”⁶ Finally, the aspectual shape is a feature which must make a difference or matter to the agent; they must be aware of the aspectual shape. The central argument of this thesis builds on Searlean intuitions about intentionality. In particular, I argue for cognitive phenomenology: a particular mode of consciousness which characterizes intentional content. The aspectual shape of a thought is the particular episode of cognitive consciousness in which the thought occurs.

b. Phenomenal Consciousness

Phenomenal consciousness is the first-personal, subjective, qualitative aspect of experience. It is most straightforwardly exhibited in perception: there is something it is like to see the color red, or to taste chocolate, which is personal to the perceiver and defies expression.

⁵ Searle targets causal-teleological views more directly in his ‘Robot Reply.’ He imagines a system whose inputs are causally connected to states of affairs in the environment. Still, these inputs would remain entirely formal symbols with no meaning to the man in the box. As a result, the causal connection fails to provide the missing element the original system lacked.

⁶ John R. Searle, “Consciousness, Unconsciousness and Intentionality,” *Philosophical Issues* 1 (1991): 52.

At a particular time, a conscious agent has a holistic phenomenology of their experience as a conscious agent. However, conscious agents can isolate particular phenomenal properties from their holistic experience. These individual phenomenal properties are called “qualia” and refer to a specific experience from one mode of sensation. For example, one can focus on the taste of an apple in isolation from the holistic experience; what it is like to taste the apple is a phenomenal property or taste-qualia.

William Lycan points out that the term ‘phenomenology’ is subject to some ambiguity. He distinguishes between ‘Q-properties,’ which are phenomenal properties of objects, and the higher-order property of what-it-is-like to experience a particular Q-property.⁷ This distinction stems from the observation that our sensory experience is ‘transparent’: when we see that an apple is red, the redness is seen as a feature of the apple, not a feature of our experience. When we focus in on what it is like to see the redness of the apple, we ascertain a higher-order property of our experience, rather than a lower-order phenomenal property of the object. For these purposes, when I write of phenomenal consciousness or phenomenal properties, I mean a property of experience or a property of the mind, rather than a property of an object. For example, a red phenomenal property is the property of a conscious agent subjectively experiencing redness, not the property of some object appearing red.

This property of the subjectivity of conscious experience was vividly brought to attention in famous papers by Thomas Nagel and Frank Jackson. Nagel argues that there is something it is like subjectively to be a bat.⁸ These what-it-is-like facts cannot be understood agent-neutrally. Because bats experience the world through echolocation, the facts about what-it-is-like to be a

⁷ William Lycan, “Phenomenal Intentionalities,” *American Philosophical Quarterly* 45, no. 3 (2008): 237.

⁸ Thomas Nagel, “What Is It Like to Be a Bat?” *The Philosophical Review* 83, no. 4 (1974): 438.

bat can only be understood if one is a bat; humans can at best imagine what it would be like to be a human living as a bat. As a result, phenomenal consciousness provides access to a class of knowledge that evades objective or third-personal description and which can only be understood subjectively.

Similarly, Jackson imagines a scientist named Mary who is raised in an entirely black-and-white environment, such that she has never seen color.⁹ Mary is a neuroscientist, and knows everything there is to know about how the brain and how it perceives color. Still, Jackson argues that Mary will learn something new upon her first perception of the color red. Mary's complete knowledge of the physical facts about color is insufficient for total knowledge about color; the information she is missing comes entirely from phenomenal consciousness.

Having isolated phenomenal properties, a question remains as to which mental states exhibit phenomenal properties. It is uncontroversial to assert that sensory experience exhibits phenomenal properties: seeing color, smelling an odor, or feeling a pain are all paradigmatic examples of phenomenal states. Some candidates for alternative phenomenal states include emotion, mood, or affect, and what Horgan and Tienson have called the phenomenology of agency. A primary project in this thesis is to argue that there exists a proprietary phenomenology of cognition. This is a mode of consciousness which represents intentional content.

⁹ Frank Jackson, "Epiphenomenal Qualia" *Philosophical Quarterly* 32 (1982): 127-136.

II. The Relationship Between Phenomenal Consciousness and Intentionality

Intentionality and Phenomenal Consciousness are two properties of the mind which have received a great deal of attention in the last fifty years. In the mid-twentieth century, the computational-functional model of the mind was a central topic of debate in the philosophy of mind. This model asserts that the phenomena of the mind can be abstracted from the physical media in which they occur. Functionalists assert that the mind is multiply realizable in any medium, so long as inputs and outputs are properly configured according to some formal program. As a result, intentionality and phenomenal consciousness were each raised independently as important properties of the mind which apparently resist functional reduction. Searle's Chinese Room argument presents a reason to doubt that the computational-functional model is sufficient for intentionality. Jackson and Nagel's arguments present reasons to doubt that the computational-functional model can account for phenomenal consciousness.

a. Separatism

Because phenomenal consciousness and intentionality were each raised independently as objections to the computational-functional model of the mind, it has not been traditionally assumed in these debates that intentionality or phenomenal consciousness bear any important relationship to one another. The position that there is no important relationship between phenomenal consciousness and intentionality is called separatism. Separatism is largely the default position in the philosophy of mind, since most of the literature on phenomenal consciousness avoids concerns of intentionality, and most of the literature on intentionality avoids concerns of phenomenal consciousness.

Formal arguments in favor of separatism are hard to come by, in part because of its status as a default view. However, one reason to think that intentionality and phenomenal consciousness are unrelated involves externalist intuitions about intentionality. Recall that it is a widely-held view that intentional content is wide: intrinsic duplicates do not necessarily share the same representational properties. However, it is also widely-held that phenomenal content is narrow: intrinsic duplicates necessarily share the same phenomenal properties. If both of these views are true, then intentionality and phenomenal consciousness cannot be strongly related to one another, since a pair of intrinsic duplicates could plausibly share all their phenomenal properties but share none of the same representational properties. Many philosophers who endorse both positions endorse separatism.

Another reason to endorse separatism is that many mental states appear to decisively exhibit either phenomenal consciousness or intentionality while decisively lacking the other. For example, standing propositional attitude states appear to exhibit intentionality while lacking any phenomenology. The belief that cats are mammals may not enter my conscious awareness very often, but I would readily endorse this belief if prompted. As a result, it makes some intuitive sense to say that I believe that cats are mammals over time as a standing propositional attitude. Still, because this state exists outside my conscious awareness, it must lack phenomenal properties. Similarly, moods may be states which clearly possess phenomenal properties, yet do not represent anything. If these types of states are possible, then the instantiation of phenomenal properties is possible without the instantiation of representational properties, and vice versa.

In recent years, separatism has been challenged by two different classes of views: representationalism and phenomenal intentionality theory. Both of these positions are

anti-separatist because they assert some strong relation between phenomenal consciousness and intentionality.

b. Representationalism

Representationalism is the thesis that every phenomenal property is identical to some representational property. According to these view, the properties of phenomenal consciousness are among the many ways by which the mind can represent content. Representationalism is thus a theory of phenomenal consciousness which grounds phenomenal content in representational content. Representationalism can be reductive or non-reductive; it simply asserts that whatever the explanatory status of representational content, phenomenal content has the same status.

However, representationalism is most commonly defended in the form of tracking representationalism, such as the theories of Fred Dretske, Michael Tye, and William Lycan.¹⁰ Tracking representationalism can be seen as an extension of the causal-teleological approach to intentionality outlined above. In addition to asserting that every phenomenal property is identical to some representational property, tracking representationalism also asserts that intentionality is explained by a tracking relation with the environment. For example, Dretske endorses a tracking relation according to which organisms develop systems with the natural function of indicating certain features of their environment. As a result, tracking representationalists are reductive with regard to representation. They assert that the property of intentionality can be explained in terms of simpler properties of natural systems. Furthermore, because phenomenal consciousness is a

¹⁰ Fred Dretske, *Naturalizing the Mind* (Oxford: Oxford University Press, 1995); Michael Tye, *Ten Problems of Consciousness* (Cambridge: MIT Press, 1995); William Lycan, *Consciousness and Experience* (Cambridge: MIT Press, 1996).

For a detailed summary and criticism of reductive representationalism, see Angela Mendelovici and David Bourget, "Tracking Representationalism," in *Philosophy of Mind: The Key Thinkers*, ed. Andrew Bailey (New York: Continuum, 2014): 209-235.

particular kind of representational property, it too can be reduced to simpler properties. Thus, tracking representationalism offers a unified, reductive and physicalist theory of mind.

One motivation for tracking representationalism is the argument from transparency. Representationalists note that when we experience qualitative properties, we do not typically experience them as properties of our experience, but rather, as properties of external objects. Though we can internally reflect upon our experience, the experience is exhausted by the properties ascribed to external objects. As a result, representationalists argue that the phenomenal properties are exhausted by the representational properties. There is no qualitative aspect of experience aside from the aspect that allows for representation of the external world. Tracking representationalism allows for a neat explanation of sensory qualities by reducing them to intentional properties, which themselves can be reduced to tracking functions.

Representationalism faces the standard anti-separatist concern: they must explain how narrow phenomenal content could be identical to wide representational content. The standard approach is to endorse phenomenal externalism, which holds that phenomenal content, like intentional content, is wide. Intrinsic duplicates do not necessarily share the same phenomenal properties, and phenomenal consciousness depends in part on factors external to the conscious agent.¹¹

In addition, representationalists require an explanation of supposed phenomenal states which lack any representational content. Because representationalists assert that every phenomenal property is identical to some representational property, phenomenal states which

¹¹ Fred Dretske, "Phenomenal Externalism, or If Meanings Ain't in the Head, Where Are Qualia?" *Philosophical Issues* 7 (1996): 143-158; William Lycan, "The Case for Phenomenal Externalism," *Philosophical Perspectives* 15 (2001): 17-35. Because representationalists take phenomenal properties to be ontologically dependent on representational properties, they assert that ordinary perception is laden with representational content such that it varies among intrinsic duplicates insofar as their intentional content also varies.

lack representational content should not be possible. Some representationalists, like Michael Tye, argue that pains in fact represent both a location in the body and some quality of danger or badness. Similarly, emotions are treated as having certain evaluative representational content about the agent's condition, and also represent certain accompanying bodily conditions (such as an elevated heart rate). Moods are slightly more difficult to explain in this way, and representationalists have offered a range of approaches to explaining what representational content such phenomenal states represent, such as representing affective properties of the world in general. Similarly, some philosophers have raised counterexamples to representationalism by considering images for which the viewer can change the phenomenal character of the viewing experience without changing any representational properties of the image.

Another difficulty for representationalism is to explain cases of perceptual distortion. These are cases in which a perceptual apparatus partially fails, such that the perception is distorted, but the experience is not regarded as misrepresenting what it aims to represent. For example, a blurry visual perception of a chair is still taken to represent a chair. As a result, a blurry perception of a chair is phenomenally different from a vivid perception of a chair, despite both perceptions having the same representational content. Some representationalists (e.g. Dretske) claim that blurry vision in fact represents blurry objects, and others (e.g. Tye) argue that blurry vision has a representational property of visual ambiguity with regard to contours which accounts for the phenomenal contrast. This deficiency in detail is taken to be a representational difference corresponding to the phenomenal difference.

c. Phenomenal Intentionality Theory

Phenomenal intentionality theory includes a broader category of views than representationalism. However, most views which fall under the general umbrella of phenomenal intentionality endorse the following theses (inspired by Uriah Kriegel's "The Phenomenal Intentionality Research Program"): (a) that there is a kind of intentionality, phenomenal intentionality, which is grounded in phenomenal consciousness; (b) that the intentional character of phenomenal intentionality is inseparable from its phenomenal character; (c) that phenomenal intentionality is narrow, in that it does not depend on anything external to the subject; (d) and that phenomenal intentionality is the basic form of intentionality from which all other intentionality is derived.

Particular theories may vary in the way they choose to spell out each thesis. For example, the grounding relation specified in (a) could include a number of different approaches. In general, phenomenal intentionality theorists endorse either an asymmetric grounding relation (such that the intentional properties of a state of phenomenal intentionality are grounded in phenomenal properties, but the phenomenal properties are not exhausted by their intentional content), or a symmetric identity relation (such that the intentional properties and the phenomenal properties of a state of phenomenal intentionality are one and the same).

Similarly, theories vary widely in their particular approach to defending the narrowness thesis outlined in (c). Some approaches are moderate, asserting that though phenomenal intentionality is constituted only by narrow phenomenal consciousness, other types of intentionality exist which are constituted by wide external conditions. However, stronger approaches assert that phenomenal intentionality is the only kind of intentionality exhibited by

mental content, and thus all possible intentionality is narrow. Because of the strength and influence of externalist intuitions, such strong approaches require a great deal of support. The primary project of this thesis is to outline a form of phenomenal intentionality theory which powerfully explains our capacity to have mental content while also avoiding the concerns of externalism.

Phenomenal intentionality theories also vary with regard to the scope of phenomenal intentionality. Some phenomenal intentionality theories see phenomenal consciousness as limited primarily to perceptual phenomenology, and aim to explain how such phenomenology is necessarily intentional. These theories can be understood as forms of non-reductive representationalism: every phenomenal property (limited to perceptual phenomenology) is identical to some representational property; however these representational properties are themselves both inseparable from their phenomenal character and explanatorily basic. In contrast, other phenomenal intentionality theories see phenomenology as far more expansive, including the phenomenologies of emotion, agency, and conscious thought. In particular, the phenomenology of conscious thought, often called the phenomenology of cognition, has garnered a great deal of attention. The form of phenomenal intentionality theory I argue for in this thesis focuses primarily on the phenomenology of cognition. Projects which focus on the intentionality exhibited by perceptual phenomenology are potentially compatible with this project. While one possible result is that all phenomenal states and all intentional states are fundamentally interdependent, it is also conceivable that the phenomenology of cognition is the only form of phenomenal consciousness which instantiates phenomenal intentionality. For now, I remain noncommittal as to these positions.

Limiting my focus now to phenomenal intentionality projects which focus on the phenomenology of cognition, there are a number of motivations for phenomenal intentionality theory. First, a number of philosophers appeal to phenomenal contrast cases. For example, David Pitt considers “multiple-center embedded” sentences, such as “the boy the man the girl saw chased fled”.¹² At first encounter, these sentences read as if they were ungrammatical strings of words. However, once one realizes that the sentence is equivalent to “the boy, who was chased by the man the girl saw, fled,” the sentence reads entirely differently. The difference between these two experiences is clearly a phenomenal difference: there is something it is like to fail to understand the sentence, and there is something it is like to understand it. The shift between these states is immediately noticeable. These aspects are unique to, and indeed define, phenomenal experiences. Since the only difference between the two states is cognitive, there must be a phenomenology which characterizes cognition.

Phenomenal contrast cases tend to work primarily as illustrative examples, rather than formal argument. A more powerful form of argument, which I explore in detail in Chapter Two, is based on the particular kind of introspective access one has to the contents of one’s own thoughts. Uriah Kriegel outlines the general form of the argument: “one has a special, immediate access to some of one’s cognitive states (and their contents); only to phenomenal states (and contents) can one have this kind of special access; therefore, (some of) one’s cognitive states (and their contents) are phenomenal.”¹³ These arguments from introspection aim to characterize the features of phenomenology-in-general, independently from any particular phenomenal mode,

¹² David Pitt, “The Phenomenology of Cognition or ‘What Is It like to Think That P?’,” *Philosophy and Phenomenological Research* 69, no. 1 (2004): 2.

¹³ Uriah Kriegel, “The Phenomenal Intentionality Research Program,” in *Phenomenal Intentionality*, ed. Uriah Kriegel (Oxford: Oxford University Press, 2013).

then argue that our access to our cognitive contents is made possible by it being a type of phenomenology-in-general. Pitt argues that in particular, our access to our cognitive contents allows us to distinguish thoughts from other mental states, distinguish thoughts from one another, and individuate the contents of our thoughts. As a result, he asserts that cognitive phenomenology is proprietary, distinctive, and individuating. This argument will be the basis from which I outline my own form of phenomenal intentionality theory.

d. Non-Reductive Representationalism

David Chalmers provides a helpful characterization of the difference between reductive representationalism and non-reductive representationalism. The key distinction is that reductive theories assert that phenomenal properties are identical to or grounded in certain representational properties “that can be understood without appeal to phenomenal notions,” whereas non-reductive theories assert that such representational properties “cannot be understood without appeal to phenomenal notions.” Some forms of PIT are reductive; these would assert that representational properties are grounded in phenomenal properties, and these phenomenal properties can be understood without appeal to representational notions. However, my version of PIT endorses an identity view and is non-reductive: it asserts that occurrent conscious thoughts are phenomenal-intentional states, possessing both representational and phenomenal properties.

Chalmers explains the nature of the relationship different theories assert between consciousness and intentionality in terms of entailment. If property r entails property p , whatever has property r must have property p . Two properties are identical if they entail each other: whatever has r must have p , and whatever has p must have r . Thus, representationalism asserts that phenomenal properties entail representational properties; a state cannot be conscious without

being intentional. In contrast, PIT asserts that representational properties entail phenomenal properties; a state cannot be intentional without being conscious. Chalmers argues that it is implausible that representational properties entail phenomenal properties because “it seems that most or all representational contents can be represented *unconsciously*, without any associated phenomenal character at all.”¹⁴ However, many PIT theories, including my own, deny the existence of any unconscious states with intentional content.

Non-reductive representationalism differs from PIT primarily in their explananda. PIT aims to provide an account of what makes a state intentional: that it has a certain kind of consciousness. As a result, PIT can plausibly assert that some phenomenal states are non-intentional, whereas others are intentional. PIT (at least in its strongest form) cannot plausibly assert, however, that some intentional states are phenomenal but others are non-phenomenal. In contrast, non-reductive representationalism aims to provide an account of what makes a state conscious: that it has intentionality. Such a theory can plausibly assert that some intentional states are non-phenomenal, whereas others are phenomenal. However, they cannot plausibly assert that some phenomenal states are intentional, whereas others are not.

My theory endorses PIT and is non-committal as far as NRR. What is essential to my theory is that no state has (full) intentional content unless it has holistic cognitive phenomenology. My theory is compatible with the view that other kinds of non-cognitive phenomenology lack intentional content. I am sympathetic, though officially non-committal, to the view that sensory phenomenology has intentional content.

¹⁴ Chalmers, “The Representational Character of Experience.”

Chapter Two: Cognitive Phenomenal Holism

I. Introduction

Some of the intuitions behind Searle's view of intentionality can motivate phenomenal intentionality theory. Searle's Chinese Room argument primarily runs as a negative argument against functionalist and causal-teleological theories of intentionality, as pointed out earlier.¹⁵ Despite the functional organization of the system or its causal connection to the world, the system lacks some condition necessary for intentionality. Searle's argument suggests that in order for a system to possess intentionality, it must possess a phenomenal awareness of its intentional content. The Chinese Room and its implications have been hotly debated for decades, and it is not meant to act as the primary argument for PIT; however, the intuitions it motivates are useful in getting a grasp on the nature of cognitive phenomenology.

His argument has two primary upshots. The first, as discussed already, is that a system can possess all the properties needed to satisfy functionalist or causal-teleological theories and still lack intentionality; therefore, such properties, even if necessary for intentionality, are insufficient. The second upshot is that such a system can behave *as if* it were fully intentional despite lacking intentionality; apparently intentional behavior is insufficient for intentionality. One implication of this second point is that from an external perspective, any apparently intentional system could conceivably lack intentionality; only from an internal perspective within the system (e.g. the first person perspective of a conscious agent) can one know for certain that a system is intentional. This can be seen as an extension of the problem of other minds from phenomenal consciousness to intentionality: just as one cannot know that a person who behaves

¹⁵ John R. Searle, "Minds, Brains, and Programs," *The Behavioral and Brain Sciences* 3, no. 3 (1980): 417-424.

as if they are in pain genuinely feels pain (has a phenomenal conscious experience), one cannot know that a person who behaves as if they represent certain content genuinely represents that content (has original intentionality).

The internal perspective provides a degree of epistemic access which is lacking in the external perspective. David Pitt's argument for cognitive phenomenology, which I review in greater detail in the following section, follows from the unique kind of epistemic access conscious agents have to their own intentional content.¹⁶ He argues that this degree of epistemic access can only be explained if there is a proprietary, distinctive, and individuating phenomenology of cognition. A theory which focuses on this phenomenology of cognition can explain the deep connection between an agent's phenomenal awareness of their intentional content and the representational properties of the content itself.

In this chapter, I outline a phenomenal intentionality theory which begins from this unique degree of internal epistemic access to intentional content. A primary desiderata for a theory of cognitive phenomenology is to account for the rich, determinate, and complex intentional contents we are capable of representing. Mendelovici and Bourget deny that intentional contents exhibit this degree of richness and complexity. They argue that in ordinary thought, our intentional contents are vague and gisty. For example, they consider a philosopher who uses the term 'supervenience.'¹⁷ When her colleague asks her to define the term, it takes her some time to arrive at a formal definition. As a result, Mendelovici and Bourget argue that her cognitive phenomenology was only sufficient to account for her initial usage of the word, which

¹⁶ David Pitt, "The Phenomenology of Cognition or 'What Is It like to Think That P?,'" *Philosophy and Phenomenological Research* 69, no. 1 (2004): 2.

¹⁷ Angela Mendelovici and David Bourget, "Consciousness and Intentionality" in *Oxford Handbook of Consciousness*, ed. Uriah Kriegel (New York: Oxford University Press, forthcoming): 30-31.

was vague and indeterminate in content. When articulating a formal definition, they argue, she merely manifested a certain behavioral disposition to describe the meaning in that way; the more determinate formal definition was never present in her occurrent cognitive phenomenology, and was only accessible by means of further inference. I am sympathetic to Pitt's argument that at least in some cases, cognitive phenomenology is determinate enough to ground rich intentional content. However, a benefit of my account of cognitive phenomenal holism is that it can account for a spectrum of determinacy of intentional content, accounting for both gisty and fully determinate intentional content.

This project of explaining the determinacy of intentional content in terms of cognitive phenomenology faces a two-fold challenge. First, it is often assumed that intentional content obeys a logical structure, yet it is not immediately clear that a phenomenology could provide such a logical structure. Many theories of mental content explain this logical structure by asserting that mental content is propositional, and that the logical structure of the proposition is isomorphic to the logical structure of the thought. Second, intentional content must be communicable. We use language to communicate thoughts to other people, and as a result, it must be possible that two individuals can share the same intentional content. Again, the propositional model is often the answer to this problem: if mental content is propositional, then so long as my thought expresses the same proposition as yours, our thoughts are about the same content. Further, because sentences in a language express propositions, this model allows mental content to be communicated between people.

Some PIT theories also rely on this propositional model of mental content to explain logical structure. For example, in his 2004 paper, Pitt argues that cognitive phenomenologies

express propositions,¹⁸ and in his 2009 paper, he turns to a type/token model according to which particular cognitive phenomenologies are tokens of phenomenal intentional types, where these types correspond to propositions.¹⁹ The problem, however, is that the propositional model contains the seeds of an argument for externalism.

The reason that the propositional model leads to externalism is precisely because of its function in allowing communication between people. If we believe that different people entertain the same intentional content because their thoughts express the same propositions, then the intuitions behind externalism can quickly arise. Burge's argument for externalism is based on the principle that, all things being equal, we should interpret people's expressions of their mental content literally.²⁰ Thus, when Bert says that he has arthritis in his thigh, we should take him to be talking about arthritis, despite the fact that arthritis is a condition of the joints. This leads Burge to argue that our mental content depends in part upon experts in our language communities who define the concepts which occur in our thoughts; as a result, our mental content depends in part upon factors external to the individual. Yet the principle of literal interpretation, from which this argument proceeds, derives directly from the propositional model. The reason to believe that Bert should be interpreted literally is that we believe that his thought about arthritis expresses the same proposition as our own; we take him to be talking (and thinking) about the same thing as us, because we believe that generally, our thoughts are all based upon the same propositions.

The account of cognitive phenomenal holism presented in this chapter aims to meet these challenges. I argue that it can account for the degree of richness, determinacy, and complexity of

¹⁸ Ibid.

¹⁹ Pitt, "Intentional Psychologism," *Philosophical Studies* 146, no. 1 (2009): 120.

²⁰ Tyler Burge, "Individualism and the Mental," *Midwest Studies in Philosophy* 4, no. 1 (1979): 73-122.

intentional content in purely phenomenological terms while rejecting the propositional model of mental content. Not only does this model allow for a spectrum of determinacy of content, ranging from vague and gisty to fully specified, it also allows for a range of similarity between the mental content of different individuals. While in some (many) cases, two people's cognitive phenomenologies may align sufficiently so as to reasonably assert that their thoughts are about the same content, it is also possible that in many cases, two cognitive phenomenologies may be so different that the two people are simply talking past each other.

II. Cognitive Phenomenology

Because my account intends to focus primarily on the nature of cognitive phenomenology, it is useful to begin by reconstructing Pitt's arguments for cognitive phenomenology.²¹ Pitt's primary argument is from introspection: (K1) One can immediately and introspectively access the contents of one's occurrent conscious thoughts; (K2) One could not possibly immediately and introspectively access the contents of one's occurrent conscious thoughts unless each thought had a proprietary, distinctive, and individuating phenomenology; therefore (K3) Every conscious thought that p , for all thinkable contents p , has a proprietary, distinctive, and individuating phenomenology.

Pitt first employs Dretske's distinction between simple seeing, the capacity to distinguish an object from its environment on the basis of its visual appearance, and epistemic seeing: "One need not be able to identify what one is seeing (know what it is) in order to be able to distinguish it from its environment."²² In contrast, epistemic seeing necessarily involves belief that one is

²¹ Pitt, "The Phenomenology of Cognition," 8-13

²² Ibid, 9.

seeing some particular object. Pitt extends this to the other modes of perceptual experience; for example, one can smell that the brownies are in the oven or feel that the surface is rough. Pitt argues that epistemic perception constitutes knowledge, and in particular, what Russell calls knowledge by acquaintance;²³ in contrast, simple perception is a form of simple acquaintance. Moreover, Pitt argues that one may have simple acquaintance with and knowledge by acquaintance of one's occurrent conscious thoughts. Not only can one introspectively distinguish a particular thought from one's other thoughts, one can also identify it as a particular thought. The former capacity is analogous to simple acquaintance, while the latter is analogous to knowledge by acquaintance. In some cases, one may not be able to classify a particular mental state (such as a fleeting sensation or emotion), but the mental state is nonetheless experienced. This is a case of introspective acquaintance without knowledge by acquaintance.

Pitt calls the particular form of introspective knowledge by acquaintance 'grasping,' where 'grasping that *p*' is analogous to 'seeing that *p*' or 'smelling that *p*.' Importantly, grasping that *p* is different from merely thinking that *p*: grasping that *p* involves having a higher order thought about *p*: "the belief that the thought *t* has the content that *p* is conclusively justified by the experience of *t*." Thus, Pitt's argument is that the immediate knowledge of one's occurrent conscious thoughts is a form of knowledge by acquaintance, and that "this is not possible without simple acquaintance, which itself depends upon the introspected state having phenomenal character."²⁴ According to Pitt, the phenomenology of cognition is a proprietary mode of consciousness just as visual and auditory perception are proprietary modes of consciousness. Within each mode, mental states can be distinguished from each other according to their

²³ Russell's theory of knowledge and its extensive commentary are not relevant here; Pitt is simply referring to knowledge on the basis of an object appearing some distinguishable way in perceptual experience.

²⁴ Ibid, 11.

phenomenal characters. Moreover, their phenomenal properties allow for knowledge by acquaintance.

Pitt finds K1, that we have immediate, introspective knowledge of our conscious occurrent thoughts, difficult to deny. He focuses on objections to K2, that such immediate knowledge can only be explained by cognitive phenomenology. He considers whether a functional/computational model of belief formation can explain this kind of access. For example, suppose that whenever the thought t consciously occurs, the thought t' co-occurs, where the content of t' is 'the content of t is p .'²⁵ Because this mechanism could occur without any phenomenology at all, it presents a counterexample to cognitive phenomenology.

However, Pitt denies that the causal co-occurrence of t and t' qualifies as *identifying* t as the thought that p so as to qualify as a form of knowledge by acquaintance. Pitt argues that it is essential that objects be identified on the basis of simple acquaintance. This kind of discrimination "requires that the object appear to one in some discriminate way, and that one be attending to it."²⁶ Thus, the mere fact that the appropriate thoughts occur in the appropriate causal arrangement is insufficient, as the higher order thought must be formed not merely at the same time as the lower order thought, but on the basis of some phenomenal features of the lower order thought. That is, the knowledge of content, in order to constitute the right kind of knowledge, must be formed on the basis of distinguishing phenomenal features which are noticed by the agent. It is not enough to have knowledge of the content of the thought; that knowledge must be essentially connected to certain phenomenal features which make the thought identifiable in consciousness in the first place

²⁵ Ibid, 19.

²⁶ Ibid.

One important feature of Pitt's view is that cognitive phenomenology is not merely sensory phenomenology which reliably accompanies the conscious occurrence of particular thoughts. Assume that every thought that 'snow is white' is accompanied by the visual phenomenology of seeing white snow. We can introspectively differentiate between our thought that snow is white and the visual phenomenology of seeing white snow; there must be something *distinctive* about the phenomenology of our thought, which is not explained by the visual phenomenology. This same line of reasoning applies to objections that thoughts are accompanied by phonological or orthographical phenomenologies: "if the tokening of a sentence in one's head is not sufficient for thinking the thought it expresses, then introspective awareness of the phenomenal properties of the sentence-image cannot be sufficient for introspective awareness of what is thought."²⁷ In general, phenomenal contrast cases provide a useful heuristic for ruling out alternative modes of phenomenology. If the alternative phenomenology is possible in the absence of one's entertaining the relevant cognitive content, then the alternative phenomenology is phenomenally distinct from the relevant cognitive phenomenology.

In his 2004 paper, Pitt asserts a relatively weak relation between phenomenology and intentionality: cognitive phenomenologies express particular intentional contents, though the two are distinguishable from each other. Later, Pitt argues for an identity view: particular cognitive phenomenologies are identical to their intentional content. A particular cognitive phenomenology is a conscious state with some particular intentional content. Pitt provides some reasons to prefer an identity relation between intentional contents and cognitive phenomenologies rather than a weaker relation. First, this view is more economical, as it does

²⁷ Ibid, 24.

not need to provide an account of independent intentional contents which would then be phenomenally constituted; rather, the phenomenology is the content. Second, this view provides a unified account of the mental. For any type of mental state, what it means to be in that state is to be having a particular kind of phenomenology. All mental states are thus uniformly phenomenally constituted. Finally, the identity view can best explain immediate introspective knowledge of mental content.²⁸ Recall that earlier, Pitt argued that we can immediately and introspectively (non-inferentially) access the contents of our thoughts, and that this introspection would be impossible unless there was a phenomenology of cognition. If the phenomenology of cognition is identical to intentional contents, the explanation can stop there. If instead the phenomenology of cognition bears some weaker relation to intentional contents (such as metaphysical grounding), the account needs a further explanation of how introspecting the phenomenology of cognition provides access to the content.

III. Cognitive Phenomenal Holism

My argument for cognitive phenomenal holism, the argument from additional mental content, appeals to Pitt's introspective argument for cognitive phenomenology. Pitt argues that the only explanation of our immediate, non-inferential access to the contents of our thoughts is that our thoughts are or are constituted by cognitive phenomenologies. Consider such a case of introspection. It is true, as Pitt suggests, that I can introspect on my thoughts and immediately know that I am thinking that snow is white. Yet further, in the act of thinking that snow is white, I also understand what 'snow is white' means. In order for such understanding to be possible, I must know certain things about snow and whiteness (snow is cold, snow falls in winter, white is

²⁸ Pitt, "Intentional Psychologism," 120.

a color, etc.). Now, just as my access to the content was immediate and non-inferential, my ability to understand the content is also immediate and non-inferential. I do not need to further reflect on snow to know that snow is cold. Insofar as it is conceptually possible for me to think about snow in the thought that snow is white, I must be thinking of something cold. Were I not thinking about a cold substance, I would not be thinking of my snow-concept. Importantly, one's access to this additional content is an extension of 'grasping' or knowledge by acquaintance. This content does not merely reliably co-occur with one's conscious thoughts; rather, one's access is based on the phenomenal features which make simple acquaintance possible. This phenomenon is important not because additional content occurs alongside one's conscious thought *t*, but that this additional content contributes to semantically determine the content of *t*.

This consideration indicates that an explanation of introspective access to belief contents must account for this 'additional content' in conscious belief. For an individual to understand the meaning of 'snow is white,' their conscious thought must include this additional intentional content. For this reason, I argue that cognitive phenomenology is holistic. The meaning of any particular belief for an individual is determined by that belief's role in a holistic Network of beliefs. Holism asserts that a change to any belief in the Network will change the meaning of every other belief in the Network; as a result, meaning can only be determined on the level of the Network, rather than on the level of particular beliefs.

Meaning holism is not a new idea to the philosophy of mind and language.²⁹ However, because this form of meaning holism forms part of a phenomenal intentionality theory, it is a particular variant of holism with certain features. First, the theory of cognitive phenomenal

²⁹ For example, W.V.O. Quine, "Two Dogmas of Empiricism," in *From a Logical Point of View* (Cambridge: Harvard University Press, 1951); Robert Brandom, *Making it Explicit* (Cambridge: Harvard University Press, 1994); Ned Block, "An Argument for Holism" *Proceedings of the Aristotelian Society* 95 (1995): 151-164.

holism is asserted as a theory of mental content, not as a theory of language-meaning. It asserts that occurrent conscious thoughts possess the contents they possess because of the role of those contents in a complex, meaning-determining Network. One implication of this view is that belief-content is never directly isomorphic to propositional content. On a simple account, this is because an individual belief's content, were it to be fully specified, would require a schematic map of every belief in the entire Network; this map both exceeds the scope of propositional content and contains relational information not captured by mere occurrence in a proposition. Yet more deeply, such a schematic map is not actually possible to make. One can imagine a diagram which lists an enormous number of propositions, connects them by lines, and clusters them into concepts. However, such an analogy relies on the possibility of isolatable atoms of belief-contents in order to construct such a diagram.

Instead, cognitive phenomenal holism occurs at the level of concepts, rather than individual beliefs. The reason the belief that snow is white contains additional content is that it is a belief about the concept 'snow', and the concept 'snow' contains a great deal of content, including content about the coldness of snow. Rather than thinking of inferential connections between atomic beliefs, it is more useful to think of overlap between concepts. What I have called the atomic belief 'snow is white' is rather a domain of overlap between my 'snow' concept and my 'whiteness' concept, which includes but is not limited to content about the whiteness of snow. However, the convention of discussing atomic beliefs to describe synchronic moments of cognitive phenomenology and to characterize the Network with a web metaphor is a useful conceptual tool, so I will continue to employ it.

To motivate this concept-level view of cognitive phenomenal holism, consider the following introspective procedure. Think of snow. Don't think anything in particular about snow, simply think of snow. Once you have a clear semantic image of snow in mind, proceed to think that snow is white. Did anything change? Most likely, simply by thinking of snow, you were already thinking that snow was white; the belief that snow is white plays an important role in contributing to form the concept snow. However, by thinking that snow is white, you shift your focus from the properties in general which determine your snow concept to one particular property. To think that snow is white is not to add any new content to your simply thinking of snow simpliciter; that content was already there, allowing you to think of snow simpliciter. This is the reverse of standard phenomenal contrast arguments: because I take cognitive phenomenology to already be on the table, this argument establishes that cognitive phenomenal holism is a concept-level phenomenon by pointing to a lack of phenomenal contrast between entertaining a concept and entertaining a proposition about that concept. Not all of your beliefs about snow will be present when merely thinking about the concept snow. In such a case, those beliefs simply are not relevant enough to your snow concept in the occurrent context so as to fall within your cognitive phenomenal threshold at the time. I explain synchronic cognitive phenomenology in the following section.

Of course, thinking that snow is white is not phenomenally identical to thinking simply of the concept snow. For one, the thought that snow is white is more determinate than the thought of the concept snow. Thinking of the concept snow is determinate in one sense: you are determinately not thinking of any of your other concepts. Yet the thought that snow is white is more fine-grained than the thought of snow itself; one difference is that the thought that snow is

white has truth conditions. I characterize the shift from thinking of snow itself to thinking that snow is white in terms of focal and peripheral phenomenology. When thinking of the concept snow, all the properties which contribute to form the concept are in consciousness to some intermediate degree: they are present to the extent that they contribute to form the concept, but none of them is explicitly in consciousness. These additional properties increase the determinacy of the content in consciousness, such that it is phenomenally distinguishable from related concepts. When shifting to think that snow is white, the particular property of snow's whiteness becomes focal to consciousness; the other properties forming the concept fade to the periphery of consciousness.

The picture of overlapping concepts is helpful to understand the phenomenology of cognitive phenomenal holism. Though the belief that snow is white carries with it additional content, this should not be understood as multiple propositional beliefs occurring simultaneously in consciousness. Rather, the additional content contributes to construct the conscious state itself; the holistic content is prior to the isolated content. The belief that snow is white is first and foremost a belief about 'snow' and 'whiteness,' and must include enough content to determinately construct those concepts. Once these concepts are constructed in consciousness, the specific content regarding the whiteness of snow can be isolated. However, this additional conceptual content is preconditional for the possibility of isolating specific content.

IV. Synchronic Cognitive Phenomenology

Given that our occurrent conscious thoughts carry with them certain additional content, it is reasonable to wonder exactly what content occurs in consciousness at a particular moment. I

refer to this moment of cognitive phenomenology as synchronic cognitive phenomenology. First, if our concepts are deeply interrelated to one another, some concepts are more closely connected than others. This property of conceptual closeness can be understood as the degree to which one concept is relevant to the meaning of another. Relevance here is a relation between beliefs which indicates the degree to which one belief constitutes the meaning of another. Given the belief that p , the more relevant a belief q is to p , the more a change to q will impact the meaning of p . Many cases suggest that relevance is a transitive property. For example, the belief that snow is white is highly relevant to the belief that snow is cold, and the belief that snow is cold is highly relevant to the belief that snow is white. Still, I am not committed to relevance being transitive in every case. There may be cases in which the belief that q is highly relevant to the belief that p , but the belief that p is only marginally relevant to the belief that q . As a result, I treat relevance as a one-directional relation. Given a conscious belief that p , the more relevant a belief q is to p , the more likely that q occurs in the synchronic cognitive phenomenology of p .

Relevance is context-sensitive, rather than a constant and abstract relation among beliefs. The circumstances during which a belief p occurs in consciousness in part determine which beliefs are relevant to p . Though generally the belief that snow falls in winter is highly relevant to the belief that snow is white, there may be contexts of occurrence of the ‘snow is white’ belief in which the belief that snow falls in winter is much less relevant. For example, if one’s thought that snow is white occurs in the context of thoughts about the Arctic, where snow falls year-round, the belief that snow falls in winter would not be as relevant. In such a case, one can plausibly be thinking of snow without thinking of a substance that falls in winter, though in ordinary cases one could not plausibly be thinking of snow without thinking of a substance that

falls in winter. Thus, relevance is a three-part relation between a conscious belief p , its context of conscious occurrence, and a belief q .

Sufficient relevance for occurrence in synchronic cognitive phenomenology depends not only on the context of the conscious occurrence of the belief, but also on the phenomenological capacity of the believer in general, and in the moment of belief. Providing necessary and sufficient conditions for sufficient relevance appears to be a daunting philosophical task. Still, the relevance relation was meant to differentiate between beliefs which must occur in synchronic cognitive phenomenology and beliefs which obviously do not occur in synchronic cognitive phenomenology. Cognitive phenomenology must have a threshold of belief content which can occur in a single synchronic moment. Characterizing this threshold will serve to clarify the ‘sufficient relevance’ clause even if it fails to provide necessary and sufficient conditions for sufficient relevance.

It is helpful to consider phenomenological thresholds in the case of sensory phenomenology. Though we are constantly bombarded with enormous quantities of perceptual input data, our sensory phenomenology includes only a small subset of that data. In general, a conscious property occurs in synchronic phenomenology if it makes a phenomenal difference: differences in perceptual input data which are not manifested in synchronic sensory phenomenology can be said to not make a difference to sensory phenomenology, just as differences in beliefs which are not manifested in synchronic cognitive phenomenology do not make a difference to cognitive phenomenology. Consider a red carpet which has an extremely subtle pattern in a slightly darker shade of red than its background. It is plausible that one could spend a large amount of time in the room with the carpet and never notice that the carpet had a

pattern; though the pattern was there all along, it did not make a difference to one's sensory phenomenology. Yet one can also focus one's attention directly on the carpet, and in so doing, notice the pattern. This suggests that one's phenomenological threshold can be expanded in cases in which we engage in deep phenomenological reflection. One might think that painters or musicians train themselves to become more sensitive to minute distinctions in sensory phenomenology, and thus expand their phenomenological thresholds.

This analogy should help clarify cognitive phenomenal threshold. Ordinarily, our synchronic cognitive phenomenology likely includes the most closely related content. We have no need to introspect more deeply; it is sufficient for meaning to have only the most relevant conceptual content in consciousness. In cases in which a belief is particularly vague or poorly formed, it may be the case that even fewer beliefs accompany it in synchronic cognitive phenomenology. The phenomenology of the belief being vague or gisty is explained by the lack of meaning-determining beliefs which would otherwise enrich the cognitive phenomenology. Furthermore, like in the sensory case, we can, in moments of deep introspection, expand our cognitive phenomenological threshold. This is most common in the case of abstract concepts. For example, conceptual analysis as a philosophical methodology appears to be precisely such an introspective project. By asking oneself "What do I mean when I think of justice," for example, one aims to reveal constitutive elements of the meaning of justice which ordinarily might escape synchronic cognitive phenomenology.

The considerations of relevance, context, and the cognitive phenomenal threshold provide the conditions for synchronic cognitive phenomenal holism. For the conscious belief that p , any other belief q is partially constitutive of the synchronic cognitive phenomenology of p if and

only if q is sufficiently relevant to p in the context of the conscious occurrence of p so as to fall within the cognitive phenomenal threshold of the believer at the synchronic moment.

This account should make clear why cognitive phenomenology can be said to be holistic without requiring one's entire network of beliefs to occur in synchronic cognitive phenomenology. If one's cognitive phenomenal threshold were sufficiently high, one could reflect upon the role of one's occurrent conscious belief within one's entire network of beliefs. This would give the believer a much more precise understanding of the meaning of their occurrent conscious belief than we ordinarily have. This implies that our ordinary beliefs are underdetermined in meaning; our synchronic cognitive phenomenologies are only as determinate as we need them to be. David Pitt makes a similar point:

Do you really know what all of the words you use mean? (What, for example, is the precise difference in meaning between, say, 'atrocious', 'execrable' and 'rebarbative'?) I think I can tell you what at least some of the words I use mean to me – which of my concepts I use them to express. I'm not so sure I could quote Webster's Third on any of them, however. (By 'atrocious', 'execrable' and 'rebarbative' I think I probably just mean very bad.)³⁰

V. Holism

Because the cognitive phenomenal threshold limits the amount of content which can occur in a moment of cognitive phenomenology, one might dispute the appropriateness of the label 'holism.' After all, cognitive phenomenology appears to be at best molecular, consisting in

³⁰ David Pitt, "Introspection, Phenomenality, and the Availability of Intentional Content" in *Cognitive Phenomenology*, ed. Tim Baynes and Michelle Montague (Oxford: Oxford University Press, 2011), 49.

closely connected segments of a much larger dispositional network of beliefs. However, even this molecular cognitive phenomenology depends on semantic holism. Our ordinary beliefs are underdetermined in meaning; our synchronic cognitive phenomenologies are only as determinate as we need them to be. Because of the limits of our cognitive phenomenal thresholds, our ordinary conscious thoughts have molecularly determined meanings. Yet in principle, if one could expand one's cognitive phenomenal threshold to a sufficient degree, one's occurrent thoughts would be semantically determined by all of one's other concepts. Expanding one's cognitive phenomenal threshold to any degree provides greater semantic determinacy of mental content. This is analogous to making a blurry visual image more clear (perhaps by entertaining the same visual object before and after putting on one's glasses). The visual or semantic object can remain constant, but one's phenomenal experience of the object can become substantially more clear or determinate in phenomenal and intentional content. As a result, the molecularism present in synchronic cognitive phenomenology depends upon the potential holism of one's Network of dispositional beliefs. Cognitive phenomenal holism is holistic for two reasons. First, the Network of dispositional beliefs is holistic. Second, any molecular cognitive phenomenology is only possible insofar as it is a selection of this holistic web of beliefs. Meanings are determined holistically, but in practice we only have partial access to those meanings.

Asserting that we only ever partially grasp the meanings of our beliefs may appear to be an unusual and undesirable implication of cognitive phenomenal holism. One consideration which may assuage this worry is that an increase in the number of meaning determining beliefs in cognitive phenomenology has diminishing returns to the degree of content determinacy. Again, an analogy with visual perception may be helpful for preliminary clarification. Once an

image reaches a certain degree of clarity, further increases to visual display capacity have only marginal effects on visual content determinacy. For example, the improvement from a standard definition television to a high definition television provides a greater increase in determinacy of visual content than a proportional increase in resolution to a high definition television. Once content reaches a suitable degree of determinacy, further increases seem to make less of a difference to phenomenology. It is reasonable to think the same is true of belief. A belief which is only supported by a handful of meaning-determining beliefs will likely be highly indeterminate. Yet once a sufficient number of meaning-determining beliefs appear in cognitive phenomenology, one's belief content reaches a suitable degree of determinacy. Further meaning-determining beliefs can certainly make the content even more determinate, but only marginally so. As a result, it should not be seen as implausible that our ordinary conscious beliefs are underdetermined in meaning relative to one's full Network.

The additional content which co-occurs alongside and semantically determines a conscious thought distinguishes concepts from each other. For example, an individual could have a relatively vague concept of precipitation, but lack concepts which distinguish snow from rain. He could know that precipitation involves the condensation of water vapor which then falls to the ground. When he comes to learn about snow and rain, his precipitation concept becomes more determinate. Furthermore, his snow and rain concepts are differentiated from each other because each contains unique additional content which makes a difference to its meaning.

VI. The Dispositional Network of Beliefs

If synchronic phenomenal holism contains only the sufficiently relevant beliefs to a particular conscious belief, then it might appear that this theory relies on the existence of

unconscious intentional states. The beliefs in the Network which do not contribute to conscious cognitive phenomenology must be states with intentional content which do not occur in consciousness. Yet if this were true, it would be at odds with strong PIT, since unconscious intentional states are states which possess intentional content but lack cognitive phenomenology. This problem can be resolved if the Network is understood as a set of dispositions to have certain cognitive phenomenologies under particular conditions. For example, when one observes one's pet cat sitting on a mat, one might be disposed to have the conscious belief, with its individuating cognitive phenomenology, that the cat is on the mat.

Pitt presents an argument against unconscious intentional states. He argues that the primary motivation for unconscious intentional states is that "there's something very wrong with saying that an individual ceases to believe what he believes when he ceases to consciously endorse it."³¹ Yet this unnaturalness can be satisfied if these unconscious states are instead understood as dispositions. It is analogously unnatural to say that an individual ceases to be a good singer when she is not singing. However, when one claims that a sleeping performer is a good singer, what one means is that, were she to sing, she would sing well. Similarly, when one claims that a sleeping believer believes that snow is white, what one means is that, were he to bring to consciousness his snow-concept, he would (consciously) believe that snow is white. In summary, "We make many acceptable ostensibly categorical claims about individuals which should be taken to be ascriptions of dispositions to be in states rather than ascriptions of being in those states."³²

³¹ David Pitt, "Conscious Belief," *Rivista Internazionale di Filosofia e Psicologia* 7 (2009): 124.

³² Ibid.

By virtue of it being a Network of dispositions, the Network is best understood as a set of non-intentional facts about the brain. For present purposes, the metaphysics of my view are not explicitly relevant. However, I do not intend my view to endorse dualism; I assume that there is some way to naturalize (though not to reduce) intentionality. In particular, I follow Searle in assuming there is something about the particular biological structure of the human brain such that it is capable of bringing about the various modes of consciousness, including cognitive phenomenology. I take it that there is some set of neurophysiological facts about an individual human's brain such that they are disposed to think about certain things in certain ways. Because these states cannot be fully understood without reference to their conscious properties, a description of these physical facts does not qualify as a reduction; rather, understanding these states would provide a causal explanation of the conscious properties.

Because at present our understanding of these neurophysiological facts is limited, it is natural to describe them in intentional language. However, because these are not cognitive phenomenal states they are not intentional states. As a result, I consider the Network, conceived of as a set of physical facts about the brain, as pre-intentional: they are a causal precondition for a human to have intentional states, but they are not themselves intentional. In contrast, physical facts about the stomach, or about other features of the brain, are fully non-intentional: not only do they lack intentionality themselves, they have no important relation to intentional states. Searle provides a helpful characterization of the binary between entirely non-intentional neurophysiological facts and the intentionality of conscious events:

“The best way to think of these matters is this: In my brain there is an enormous and complex mass of neurons embedded in glial cells. Sometimes the behavior of these

elements of this complex mass causes conscious states, including those conscious states that are parts of human actions. The conscious states have all of the color and variety that constitute our waking lives. But at the level of the mental, those are all the facts. What goes on in the brain, other than consciousness, has an occurrent reality that is neurophysiological rather than psychological. When we speak of unconscious states, we are speaking of the capacities of the brain to generate consciousness. Furthermore, some of the capacities of the brain do not generate consciousness, but rather function to fix the application of the conscious states. They enable me to walk, run, write, speak, etc.”³³

Searle’s theory of intentionality has been chronically misinterpreted as a moderate form of phenomenal intentionality theory. For example, Mendelovici and Bourget claim that “Searle (1990, 1991, 1992) takes at least some standing states, such as non-occurrent beliefs and desires, to have non-phenomenal intentionality.”³⁴ At times, Searle’s theory certainly appears to ascribe intentionality to non-occurrent states. Searle most directly addresses the issue in his 1991 paper “Consciousness, Unconsciousness, and Intentionality.” In it, he argues that unconscious mental states “must be at least possible candidates for consciousness.”³⁵ It appears that there exist a class of intentional state which is not conscious but merely potentially conscious. However, Searle later rules out this interpretation. First, he classes supposed unconscious intentional states as dispositions to have conscious thoughts. Second, he denies the existence of the unconscious altogether:

³³ John R. Searle, *The Rediscovery of the Mind* (Cambridge: MIT Press, 1992), 187-188.

³⁴ David Bourget and Angela Mendelovici, “Phenomenal Intentionality” in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta (Stanford: The Metaphysics Research Lab, 2017).

³⁵ John R. Searle, “Consciousness, Unconsciousness, and Intentionality,” *Philosophical Issues* 1 (1991): 51.

“There is nothing going on in my brain but neurophysiological processes. Those processes are capable of generating conscious states.... But of the unconscious neurophysiological features, some are mental and some are not. The difference is not in consciousness, for they are both, by hypothesis, unconscious. The difference is that the mental ones are candidates for consciousness. That's all. There isn't any aspectual shape at the level of neurons and synapses.”³⁶

Thus, what Searle calls the unconscious is to be understood as non-intentional. His point is that of the many functions of the brain, a particular set is capable of bringing about consciousness, and this set should be understood as importantly different from the functions of the brain which are unrelated to consciousness. Finally, intentionality can only be ascribed to conscious events. Searle's view is philosophically closer to strong phenomenal intentionality theories like Pitt's and my own than to moderate intentionality theories, and his work should be seen as foundational to strong PIT theories.

The account of cognitive phenomenal holism is a theory of mental content based on the argument for cognitive phenomenology presented by David Pitt. My contribution, that knowledge by acquaintance of mental content is possible only if mental content is holistic, provides the basis for a detailed description of cognitive phenomenology. For such an account to be complete, however, it must present a response to externalism about mental content. I provide such a response in the following chapter.

³⁶ Ibid, 59.

Chapter Three: The Externalist Threat

I. Externalism and its Motivations

Externalism about mental content (hereafter simply ‘externalism’), as discussed in the Introduction, is the view that factors in the external world play an essential role in determining mental content. To put it another way, externalism is the view that mental content does not entirely depend on internal mental states. It should be clear that externalism about mental content is inconsistent with PIT, including and especially cognitive phenomenal holism. Under most PIT views, including my own, mental content is determined entirely by cognitive phenomenology. Because cognitive phenomenology itself depends upon only internal mental states, PIT views assert that mental content is determined internally, independent of any external factors. Because externalism about mental content is a dominant view in the philosophy of mind, no PIT account is complete without a response to externalism.

Most existing PIT theories fail to sufficiently distance themselves from a propositional model of content, and as a result their responses to externalism are either insufficient or unmotivated and *ad hoc*. Cognitive phenomenal holism provides the resources and motivation for a private concepts response to externalism. Moreover, externalism about mental content is an incredibly unintuitive position, and is only plausible if it is the best available model for ascriptions of mental content. If a theory can solve the problems externalism intended to solve while remaining internalist about mental content, such a theory is more plausible. I offer here such an internalist theory of mental content, the private concepts theory.

There are a wide variety of positions and theories which fall under the umbrella of externalism, but the most relevant to the present debate is Tyler Burge’s social externalism,

which he first introduced in the 1979 paper “Individualism and the Mental.” Burge’s argument is the preeminent externalist threat to PIT. Burge’s argument is complex: at its core is a fairly straightforward thought experiment, his conclusions rely upon certain assumptions and intuitions. This discussion aims to articulate and evaluate those assumptions and intuitions.

To begin, Burge subscribes to a view according to which intentional contents are those which occur in noun phrases within an attitude. For example, in the sentence “Joe believes that cats are mammals,” his mental state, or attitude, is believing, and his intentional content occurs in the noun phrase “cats are mammals.” Burge calls such that-clauses ‘content-clauses’; elsewhere, I have referred to this as the propositional model of mental content. Burge asserts that “Although the notion of content is, for present purposes, ontologically neutral, I do think of it as holding a place in a systematic *theory* of mentalistic language.”³⁷ This theory, among other motivations, aims to place restrictions on when mental contents are the same or different from one another. For one, Burge expects most theories to assert that “Al intends to climb Mt. McKinley” and “Al intends to climb the tallest mountain in the United States” to represent different contents, though each thought has the same extension. Further, he asserts that it is uncontroversial that “On any systematic theory, differences in the *extension*—the actual denotation, referent, or application—of counterpart expressions in that- clauses will be semantically represented, and will, in our terms, make for differences in content.”³⁸ This theoretical foundation provides a test for when content is different: “if the content of one thought can be specified by a particular ‘that’-clause but the content of another thought cannot be specified by the same ‘that’-clause, then the contents are distinct.”³⁹

³⁷ Tyler Burge, “Individualism and the Mental,” *Midwest Studies in Philosophy* 4, no. 1 (1979): 74.

³⁸ *Ibid.*, 75.

³⁹ Juhani Yli-Vakkuri and John Hawthorne, *Narrow Content* (Oxford: Oxford University Press, 2018), 8.

With this foundation and test in place, Burge introduces a progressive thought experiment. The first scenario involves a man, Bert, who “has a large number of attitudes commonly attributed with content clauses containing ‘arthritis’ in oblique occurrence.” For example, he believes that arthritis is a disease, that his Grandma has arthritis, that arthritis is painful, etc; “In addition to these unsurprising attitudes, he thinks falsely that he has developed arthritis in the thigh.”⁴⁰ Next, Burge imagines a counterfactual situation. Ernie is an intrinsic duplicate of Bert: he is physiologically identical, down to the molecular level, and possesses the same mental states as Bert at the time that Bert thinks he has arthritis in his thigh. As Burge explains, both Bert and Ernie were disposed to believe that arthritis can occur in the thigh, “never hearing anything to prejudice him for or against applying it in the way that he does.” As a result, “The *counterfactuality* in the supposition touches only the patient’s social environment . . . in our imagined case, physicians, lexicographers, and informed laymen apply ‘arthritis’ not only to arthritis but to various other rheumatoid ailments. The standard use of the term is to be conceived to encompass the patient’s actual misuse.”⁴¹ Thus, though Ernie is also disposed to say to his doctor “I believe I have arthritis in my thigh,” in Ernie’s case, the doctor accepts his assertion as unproblematic.

The force of Burge’s argument comes from comparing the actual situation (Bert) with the counterfactual situation (Ernie). He concludes that “In the counterfactual situation, the patient [Ernie] lacks some—probably *all*—of the attitudes commonly attributed with content clauses containing ‘arthritis’ in oblique occurrence.”⁴² All of Ernie’s alleged arthritis-thoughts are in fact

⁴⁰ Burge, 77.

⁴¹ Ibid, 78.

⁴² Ibid.

thoughts about some other concept (perhaps ‘tharthritis’). This relies on the test for content outlined above: Bert’s and Ernie’s thoughts cannot be expressed by the same content clauses, because Bert’s thoughts are about arthritis whereas Ernie’s are about tharthritis. Burge concludes, “ the patient’s mental contents differ, while his entire physical and non-intentional mental histories, considered in isolation from their social context, remain the same . . . The difference in his mental contents is attributable to differences in his social environment.”⁴³ (Burge, 106). In summary, in the actual context Bert believes he has arthritis in his thigh, but in the counterfactual context Ernie does not believe he has arthritis in his thigh. Because Bert and Ernie are intrinsic duplicates, there is no internal difference between the actual and counterfactual contexts. Therefore, the difference in mental content is not explained by any internal differences. The only relevant differences between the contexts are the practices of the relevant linguistic communities. Therefore, one’s social and linguistic context determines the contents of one’s thoughts.

Burge thinks this form of argument can be extended to any kind of content: “We could have used an artifact term, an ordinary natural kind word, a color adjective, a social role term, a term for a historical style, an abstract noun, an action verb, a physical movement verb, or any of various other sorts of words.”⁴⁴ He notes that the thought experiment depends only on the possibility that someone can understand a concept only partially. He adds, however, that most cases will appear idiosyncratic since “Common linguistic errors, if entrenched, tend to become common usage.” Still, any term which is defined by experts (such as scientists) will potentially generate a Burge-case. Burge includes a series of examples, including sofas, brisket, legal

⁴³ Ibid, 79.

⁴⁴ Ibid.

contracts, and color terms. For all of these cases, it is plausible that the subject be able to correctly deploy the term in a wide range of cases, while still possessing some partial conceptual misunderstanding. Burge's argument could then be summarized as follows: partial possession of a concept is possible; concepts which are partially possessable can generate Burge-cases (in which content varies between actual and counterfactual cases); Burge-cases demonstrate that mental content depends on factors in one's linguistic community. The reasoning from partial possession of a concept to the externalist conclusion is straightforward and uncontroversial. As a result, much of the debate to follow bears on whether it is possible to possess a concept partially.

Notably, Burge recognizes that "the thought experiment does appear to depend on the possibility of someone's having a propositional attitude despite an incomplete mastery of some notion in its content," and that in contrast, "An ordinary empirical error appears not to be sufficient."⁴⁵ Burge asserts that there are a wide range of concepts which we are able to deploy in a number of everyday situations, while remaining agnostic as to certain propositions which are conceptually impossible. One further upshot Burge identifies is that social factors determine mental content even in cases where content is well-understood, "For if the social environment had been appropriately different, the contents of those attitudes would have been different."⁴⁶

Burge motivates the possibility of partial concept possession with further argument. To start, Burge says that "I assume that a primary way of achieving theoretical understanding is to concentrate on our *discourse* about mentalistic notions." In support of this point, Burge asserts that "there is a methodological bias in favor of taking natural discourse literally, other things being equal. For example, unless there are clear reasons for construing discourse as ambiguous,

⁴⁵ Ibid, 83.

⁴⁶ Ibid, 85.

elliptical, or involving special idioms, we should not so construe it. Literal interpretation is *ceteris paribus* preferred.”⁴⁷ Burge aptly recognizes, then, that the issue at stake “is whether the steps should be taken in the literal way in which I have taken them, and thus whether the conclusion I have drawn from those steps is justified.”⁴⁸ In his response to Burge, outlined later in this chapter, Pitt calls this the Principle of Literal Interpretation (PLI), and makes it his primary target.

Burge first notes that we often distinguish certain types of cases as genuine misunderstanding. His examples include a foreigner who uses an incorrect word, tongue slips and Spoonerisms, and malapropisms. For each of these, Burge argues that there are relevant conditions such that listeners can easily replace the word with its obvious substitute, given the context. Finally, particularly “radical misunderstandings” may warrant reinterpretation: “If a generally competent and reasonable speaker thinks that ‘orangutan’ applies to a fruit drink, we would be reluctant, and it would unquestionably be misleading, to take his words as revealing that he thinks he has been drinking orangutans for breakfast for the last few weeks.”⁴⁹ As a result, Burge argues that in all ordinary situations, in which none of these conditions obtain, such that the *ceteris paribus* clause applies, literal interpretation is the natural way to ascribe mental content.

Finally, Burge defends the PLI against four arguments for non-literal interpretation, the latter two of which are relevant here and which “are often invoked in tandem.” The first such argument is that for each misunderstood concept, there is some other concept the subject had in mind, “thus replacing contents that are apparently false on account of the misconception, by true

⁴⁷ Ibid, 87-88.

⁴⁸ Ibid.

⁴⁹ Ibid, 90-91.

contents.”⁵⁰ For example, in the actual scenario, Bert had in mind ‘tharthritis’, rather than ‘arthritis.’ The second method is to “to count the error of the subject as purely metalinguistic. Thus the patient’s apparent belief that he had arthritis in the thigh would be reinterpreted as a belief that ‘arthritis’ applied to something (or some disease) in his thigh.” These methods can work together, as one could claim that Bert’s thought was both about some concept ‘tharthritis,’ and that he falsely believed that this concept was called ‘arthritis.’ Burge argues that such an interpretation of the thought experiment is implausible due to its violation of our ordinary practices of mentalistic attribution. Burge asserts that “In counting beliefs as shared, we do not require, in every case, that the subjects ‘fully understand’ the notions in those belief contents, or understand them in just the same way.”⁵¹ That is, even upon learning that a subject only partially grasps a concept, we continue to ascribe to them (conceptually misinformed) beliefs about that very same concept. One reason Burge thinks that we do not reinterpret beliefs to be about a new concept, such as tharthritis, is that:

“The appropriate restrictions on the application of this term and of the patient’s supposed notion are unclear. Is just any problem in the thigh that the patient wants to call ‘arthritis’ to count as tharthritis? Are other ailments covered? What would decide? The problem is that there are no recognized standards governing the application of the new term. In such cases, the method is patently *ad hoc*.”⁵²

⁵⁰ Ibid, 93.

⁵¹ Ibid, 94.

⁵² Ibid.

Furthermore, Burge argues that such cases require us to “proliferate terminology without evident theoretical reward. We do not engender better understanding of the patient by inventing a new word and saying that he thought (correctly) that tharthritis can occur outside joints.”⁵³

As to the metalinguistic argument, Burge again argues that it “has no prima facie support as an account of the language of mentalistic ascriptions.” This is because such a view would require, according to Burge, that all of the subjects’ contentful assertions about arthritis were rather merely assertions about the term ‘arthritis.’ For example, Bert believed that his grandmother had ‘arthritis,’ that ‘arthritis’ is painful, etc. Burge repeatedly asserts that one benefit of the PLI is that it allows for an easy account of shared mental contents. Under such an account, the expert and the layperson can share beliefs about the same concepts, despite the fact that they understand the concepts to widely varying degrees. In both cases, the basic point is that “there appears to be a general presumption that a person is reasoning at the object level, other things being equal. The basis for this presumption is that metalinguistic reasoning requires a certain self-consciousness about one’s words and social institutions.”⁵⁴ As a result, Burge thinks that in all non-deviant cases (i.e. excluding Malapropisms, Spoonerisms, etc.), the PLI is the most natural method of mentalistic attribution.

II. PIT Responses to Externalism

Most PIT theorists acknowledge that externalism poses a serious threat. Though they have offered a range of responses to externalism, their strategies are only successful if they

⁵³ Ibid.

⁵⁴ Ibid, 97.

sufficiently differentiate mental content from the propositional model, truth conditions, and reference. Here, I outline a range of responses from PIT against externalism.

Some PIT theorists attempt to accommodate both the internalist intuitions of PIT and externalist theories of content. Horgan and Tienson make a weaker claim involving phenomenal intentionality: “There is a kind of intentionality, pervasive in human mental life, that is constitutively determined by phenomenology alone.”⁵⁵ This kind of intentionality is not the only kind of intentionality under their view; their claim is merely that this kind of intentionality exists. They deny strong externalist theories of intentionality: such theories “are fundamentally mistaken, because they claim to naturalize the entire phenomenon of mental intentionality and yet there is a rich and pervasive kind of narrow intentionality—phenomenal intentionality—that is constitutively independent of external factors.”⁵⁶ The authors discuss Twin Earth cases, in which truth conditions are dependent on different external objects and thus vary among phenomenal duplicates. They explain: “Referring to something, mentally or linguistically, requires appropriate relations to that thing; but having thoughts that are intentionally directed toward such a thing—thoughts purporting to refer to such a thing—does not.”⁵⁷ As a result, they establish a distinction between phenomenally identical *narrow truth conditions* and externally dependent *wide truth conditions*. As a result, “there are wide truth conditions for these thoughts that are partially determined by features in the environment that may be unknown to the thinker. But again: these wide truth conditions, differing as they do for your thoughts and your Twin

⁵⁵ Terence Horgan and John Tienson, “The Intentionality of Phenomenology and the Phenomenology of Intentionality,” in *Philosophy of Mind: Contemporary and Classical Readings*, ed. David J. Chalmers (Oxford: Oxford University Press, 2002).

⁵⁶ Ibid.

⁵⁷ Ibid.

Earth doppelgängers phenomenally identical thoughts, are grounded on shared narrow truth conditions.”⁵⁸

One problem with this view is that it leaves the status of intentionality uncertain. Why is it that intentionality is sometimes internal and sometimes external? Which of these kinds of intentionality is more fundamental to mental content? This line of response appears to make theoretical sacrifices for no reason other than to accommodate externalist intuitions; as a result, it lacks sufficient independent motivation.

Katalin Farkas makes a more ambitious claim about mental content: “if one is convinced that narrow phenomenal intentionality is legitimate, there is nothing stopping one from claiming that all intentionality is narrow.”⁵⁹ She notes that the reference of every thought, as Horgan and Tienson establish, depends in part on factors external to the subject. She calls these the state of the world or the circumstances of evaluation. For us to determine the content of a thought, we must know the external conditions of the world relative to which we can evaluate its content. Under her account, the circumstances of evaluation are a contextual feature. Importantly, that means that these features are external to the content of the thought. Because these features are what lead to the varying truth values, “The proposal is that *we treat all other features of the context as also external to the content.*” For example, in the case where Farkas thinks of Farkas’ mother and Twin Farkas thinks of Twin Farkas’ mother, “I suggest that we regard this case as two thoughts having the *same content and different object*, because of some difference *external* to the content of the thought.”⁶⁰ One implication is that thought content is independent from truth conditions; truth conditions themselves, because they are object-dependent, must be taken as

⁵⁸ Ibid.

⁵⁹ Katalin Farkas, “Phenomenal Intentionality Without Compromise,” *The Monist* 91 no. 2 (2008), 274.

⁶⁰ Ibid, 282.

features of the context and not as part of the content of the thought. Essentially, Farkas asserts that all content is narrow, whereas Horgan and Tienson assert that some content is narrow and some is wide. Farkas' crucial move is to remove reference-determining features from mental content and instead make them broadly contextual features.

Notably, unlike Horgan and Tienson, Farkas notes that "If [mental] contents are propositions traditionally conceived, then my proposal is flawed. However, in recent years there has been growing awareness that it is not compulsory to think about contents in this way."⁶¹ This is because Farkas wants content to vary across worlds and across contexts. While propositions vary in meaning from world to world, they have constant meanings across contexts. An approach like Farkas' makes sense if one has a non-propositional model of mental content. However, the anti-externalist argument ought to follow from a defense of a non-propositional model. By simply assuming that thoughts are non-propositional, without offering a defense or an alternative model of content, Farkas' argument appears *ad hoc* despite its promising approach.

Brian Loar takes a third approach to this problem. He begins from the premise that "While the internalist intuition appears to me correct, the core of current externalist theory also appears correct. So the core of externalist theory must be compatible with intentionality's being an internally constituted feature of mental states."⁶² He aims to reconcile these two positions by asserting that "Externalists are right about the reference and truth-conditions of thoughts. But despite vivid appearances to the contrary, intentionality does not presuppose reference and it is not externally determined."⁶³ Like Farkas, he aims to avoid the concerns of externalism by

⁶¹ Ibid, 283.

⁶² Brian Loar, "Phenomenal Intentionality as the Basis of Mental Content" in *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, ed. Martin Hahn and B. Ramberg (Cambridge: MIT Press, 2003), 229.

⁶³ Ibid, 233.

asserting that reference and truth conditions, which drive externalist intuitions, are not features of mental content. Loar asserts that “Mental content is in fact individuated independently of that-clauses.”⁶⁴ This denial of the propositional model is a necessary starting point because it seems straightforwardly true that propositional attitude clauses capture reference, whereas he aims to deny that mental content presupposes reference.

Still, in order to fully refute externalism, Loar must show that without the propositional model, mental content does not involve reference. He claims that the starting premise of externalism is that “thoughts cannot purport to refer unless they impose success-conditions, or satisfaction-conditions; and these depend, however indirectly, on reference to objective properties.”⁶⁵ That is, intentionality presupposes reference. The second premise, according to Loar, is that reference is determined by external relations between concepts and the properties to which they purport to refer (if the concepts are outwardly directed). In Burge’s externalism, these relations include facts about the subject’s linguistic community and its experts. As a result, intentionality presupposes relations to external conditions, so externalism is true. While Loar accepts the second premise about the conditions for reference, he denies that intentionality presupposes reference.

Loar notes that he takes the opposite approach of internalist theories of intentionality based on conceptual roles. Those theories accept the first premise (that intentionality presupposes reference) but deny the second premise (that reference is determined by external relations). However, he finds that such theories must aim to explain all apparently socially deferential concepts (like arthritis) in terms of narrow concepts, such that water is taken “to mean

⁶⁴ Ibid, 229.

⁶⁵ Ibid, 234.

something like ‘the so and so stuff that causally grounds our use of “water”’. The reference of the concept ‘water’ then varies with the reference of ‘our’.”⁶⁶ This approach must work for all concepts. In particular, he finds that spatial concepts “appear to have a crucially demonstrative element, pointing visually and tactually to certain relations and properties, at least vaguely,” and that “By pointing to the sorts of relation and properties that are to count as curviness, betweenness etc, spatial perception apparently gives worldly content to otherwise purely abstract concepts.”⁶⁷ As a result, a ‘conservative’ internalist approach lacks the narrow conceptual resources to account for all wide content.

Instead, Loar argues that a suitable toolbox of conceptual resources can be built on the basis of phenomenal intentionality. Wide concepts will depend on narrow concepts, which themselves are derived from phenomenal experience. For example, “the special directedness of recognitional concepts, including spatial and basic-object concepts, derives from the singular directedness of perceptual experience.”⁶⁸ Loar calls this apparent demonstrative property of intentional content “intentional directedness.” Loar proceeds to outline a taxonomy of concepts such that even highly socially deferential concepts like arthritis ultimately have their basis in narrow phenomenal concepts. This allows him to create a sharp distinction between intentional directedness and reference: “Intentional directedness is an object-independent property, and it does not involve relations to objects. Reference comprises various causal and other relations to objects, and it is absurd to think of those reference relations as somehow instantiated without objects.”⁶⁹ Because all the concepts necessary for abstract mental content can be formed on the

⁶⁶ Ibid, 239.

⁶⁷ Ibid, 240.

⁶⁸ Ibid, 255.

⁶⁹ Ibid, 257.

basis of intentional directedness, rather than reference, Loar finds it implausible that intentionality presupposes reference. Because this first premise was essential to externalist arguments, Loar denies that mental content itself depends on any external factors.

III. Pitt's Response

Though these PIT responses offer some promising insights, David Pitt's response is the most convincing and most compatible with the resources of cognitive phenomenal holism. Rather than constructing and motivating an alternative model for ascription of mental content, Pitt attacks Burge's argument directly, aiming to undermine the plausibility of literal ascriptions of mental content. Burge's argument is meant to elicit intuitions about what is natural to say about the mental content of individuals in certain situations. In general, Burge ascribes mental content by taking utterances to be literally true descriptions of mental content. Pitt aims to reject Burge's argument by demonstrating that it is incoherent to take many of these intuitively natural ascriptions as literally true. Burge's argument in "Individualism and the Mental" is only sound if the patient, Bert, literally believes that he has arthritis; if this belief is not literally true than Burge's argument fails to demonstrate that two identical individuals could have differing mental content due to changes in external conditions.

Recall that Burge's model is guided by the Principle of Literal Interpretation, which states that, all things being equal, natural ascriptions of mental content should be taken literally. Pitt's first questions why Bert's conceptual confusion about arthritis is not a sign that all things are not equal. Yet the power of Burge's argument is to show how natural it is to ascribe other thoughts about arthritis to Bert, despite his conceptual confusion. This is the "domineering role"

of the PLI: “The fact that linguistic meaning trumps individualistic factors in cases of [conceptual] dissonance between what a speaker believes about his words and what they mean in his language shows why there is normally a match between linguistic content and mental content.”⁷⁰ In Burge’s cases, though there is conceptual confusion, it is still intuitively natural to ascribe to them the literal content of their utterance. Pitt calls these intuitions the Burgean Intuitions: “in spite of the fact that their utterances are conceptually dissonant, it still seems perfectly appropriate to use their words in characterizing what they think.”⁷¹

Pitt concedes that these intuitions are powerful; however, he argues that other considerations override these intuitions and demand non-literal ascription of mental content from utterances. Pitt points out that Burgean cases only work in cases where the subject’s error is a conceptual error, rather than merely an empirical error. As a result, he argues that literal belief ascriptions violate the principle of charity: that “One should not, *ceteris paribus*, attribute contradictory or incoherent beliefs.”⁷² In Bert’s case, Burgean principles would ascribe to him the literal belief that he has arthritis in his thigh, a belief which is, according to Pitt, conceptually incoherent. Yet Pitt points out that it is not conceptual incoherence but “conceptual *dissonance* that is doing the work.”⁷³ As a result, he considers “dissonant-but-coherent” utterances.

Pitt considers malaprops: cases where a subject makes an utterance that is not incoherent, but implausible. For example, Pitt considers a subject who asserts that “Pat Boone was an early American frontiersman.”⁷⁴ Though it is not incoherent that she could have such a belief, the more obvious explanation is that though she said something about Pat Boone, she intended to say

⁷⁰ David Pitt, “The Burgean Intuitions,” unpublished, 12.

⁷¹ Ibid, 16.

⁷² Ibid, 17.

⁷³ Ibid, 23.

⁷⁴ Ibid, 24.

something about Daniel Boone. Though it may seem natural to say that she literally believed that Pat Boone was a frontiersman, “This really is not what she thought. . . . She simply misexpressed herself. She had a false belief about which word to use to say what she was thinking, and that belief led to her anomalous utterance.”⁷⁵ Burge concedes that malaprops are evidence that one should not subscribe content by the principle of literal interpretation in the subject’s case. Yet Pitt’s point is that the Burgean Intuitions are equally strong in the case of the malaprops. Thus, “the presence of the Intuition provides no assurance that we should apply the Principle in spite of a speaker’s errors.” This objection is more powerful than the charity objection, as it demonstrates that the motivation behind Burge’s Principle is insufficient evidence to support his model of belief ascription.

In response to Burge, Pitt asserts the Incompetence Principle: “If a speaker lacks competence with a term *t*, then his *t*-involving utterances should not be taken at face value.”⁷⁶ Pitt argues that this principle should be obvious: “If you don’t know the meaning of the words you utter, then you don’t mean what they mean, and you’re not thinking what they mean.” Pitt argues that the Incompetence Principle along with the failure of the Burgean Intuitions to guarantee literal interpretation “conspire to short circuit *any* thought experiment along Burgean lines.”⁷⁷ As a result, Pitt thinks his argument refutes all forms of Burgean externalism.

IV. Private Concepts

Pitt’s internalist rejection of externalism is internally coherent, but has far-reaching implications for mental content. Furthermore, it is not clear that Pitt’s own PIT has the resources

⁷⁵ Ibid, 26.

⁷⁶ Ibid, 30.

⁷⁷ Ibid, 32.

to support the radically private view of mental content required by his rejection of externalism. In Pitt (2009), he endorses a type/token model of cognitive phenomenology, by which individual cognitive phenomenal experiences are tokens of cognitive phenomenal types, construed as independent abstract objects. As a result, different people can share cognitive content when they token the same phenomenal type. Yet under this view, concepts are objective, making it appear to succumb to an externalist argument. For example, one might think that Bert and his doctor token the same cognitive phenomenal type with regard to arthritis, despite Bert's conceptual confusion. If instead Pitt were to assert that Bert indeed had some ambiguous 'tharthritis' concept, and that concepts can be radically private, then there would need to be a phenomenal type for each possible radically private concept. Such an account would sacrifice much of the parsimony that the type/token account initially provided.

In his more recent work, Pitt has embraced the radical privacy implied by his response to externalism. He says "We often, perhaps typically, don't think what we say, and don't say what we think; nor do we often agree in what we think, even if we agree on what words we ought to use to express it."⁷⁸ Moreover, the problem of other minds seriously applies to attributions of mental content: "Though what you say you think provides me with evidence of what you do think, the possibility that the way I construe the evidence – the content I assign to your utterance – is not the content of your thought is a live one."⁷⁹ Finally, we probably should not expect significant intersubjective similarity of concepts: "We ought not to assume that an individual's apparent competence with a given set of syntactic-phonological types is an indication that he has

⁷⁸ Pitt, "Introspection, Phenomenality, and the Availability of Intentional Content," 51.

⁷⁹ Ibid.

internalized the same unique, determinate system of linguistic meaning as anyone else who is apparently competent with that same set of syntactic-phonological types.”⁸⁰

If we must choose between a private concept view and a type/token view, it seems clear to me that the private concept view has better independent support. The development of the private concept view from natural internalist intuitions about ascription of mental content ground it as the basis for a phenomenal theory of mental content. However, the private concept view needs to be situated within a more robust phenomenal intentionality theory. I argue that my theory of cognitive phenomenal holism, for which I have already provided independent support, is consistent with a private concept view, and thus provides both a complete account of phenomenal intentionality along with a powerful response to externalism.

Consider again the conditions for synchronic cognitive phenomenal holism: For the conscious belief that *p*, any other belief *q* is partially constitutive of the synchronic cognitive phenomenology of *p* if and only if *q* is sufficiently relevant to *p* in the context of the conscious occurrence of *p* so as to fall within the cognitive phenomenal threshold of the believer at the synchronic moment. When we speak of Bert’s concept of arthritis, all that we can literally describe is a single moment of Bert’s cognitive phenomenology. Of course, we can generalize his concept in terms of his dispositions or the content which tends to occur in moments of his cognitive phenomenology. However, it is important to remember Searle’s lesson: mental content is nothing beyond conscious events and the neurophysiological facts which make conscious and nonconscious events possible; intentionality never occurs in anything other than conscious events. As a result, for two people to literally share the same content, they would need to share

⁸⁰ Ibid.

the same synchronic cognitive phenomenology at the same moment. While in principle this is possible, due to the massive number of factors determining cognitive phenomenology at any moment, it is unlikely that this ever actually occurs.

As Pitt suggests, this should not be particularly surprising. It seems likely that different people conceptualize the same public concepts in significantly different ways, especially as those concepts increase in abstraction. The private concepts theory asserts that those differences should count as genuine differences in mental content. Furthermore, the theory of cognitive phenomenal holism provides the resources to make sense of these differences. During an event of cognitive phenomenology, one's consciousness is focused on a concept, for example snow. Recall that in order for it to be possible for one to determinately focus one's consciousness on snow, rather than a related concept (such as rain), one's cognitive phenomenology must include content which distinguishes between these two concepts (such that one can know by acquaintance that one is thinking of snow). For example, when one thinks that snow is white, one must implicitly have in one's cognitive phenomenology that snow is cold, that snow falls in winter, etc. These related concepts occur in the background of consciousness insofar as their presence makes a difference to the occurrent cognitive phenomenology. Though one is not focally aware of this content, it makes determinate the content which is one's focal conscious awareness. As a result, one's private concept is simply whichever cluster of related concepts occur in one's background consciousness and ground one's cognitive phenomenology.

One difficulty of this account is that cognitive phenomenology varies from moment to moment; as a result, private concepts themselves would vary too often to be useful descriptions of mental content. There are two helpful ways to generalize an individual's private concepts.

First, it is useful to think of cognitive phenomenology as a conscious event, rather than as discrete moments. There is precedent for this in perceptual phenomenology: for example, when I burn my hand, I at first feel a cold sensation, then a steady increase in pain culminating in an intense burn, which gradually dulls until I can no longer feel it. Though this description includes a fluctuating range of discrete phenomenologies, I can refer to the totality as the conscious experience of burning my hand. In the case of cognitive phenomenology, this could include the conscious experience of thinking about the whiteness of snow. Though one may entertain a range of related concepts over the course of this event, we can speak generally of the event as a whole.

Second, when comparing private concepts over longer time periods, it may be useful to speak in terms of one's dispositions to have certain events of cognitive phenomenology. Again, these should not be thought of as 'standing beliefs' or 'standing propositional attitudes' which have intentional status. As Searle emphasizes, conscious events are the only things which can be ascribed intentionality; there is no 'unconscious' beyond the neurophysiological facts about the brain, and those facts are not candidates for intentionality. This should not be seen as a return to a type/token model. One's particular cognitive phenomenologies are not tokens of an objective cognitive/phenomenal type. Rather, they are manifestations of one's dispositions to have certain phenomenologies in particular circumstances. They are individually personalized, dependent on the individual's neuroanatomy, and highly context-sensitive generalizations.

Finally, when describing a private concept it is important to keep in mind that not every concept which occurs in one's cognitive phenomenology plays an equal role in determining the meaning of the cognitive phenomenal event. Recall the spectrum of relevancy: some related concepts and properties are more relevant, and thus play a larger meaning-determining role, than

other concepts. As a result, a description of a private concept should give priority to the core conceptual features of the concept. For example, when comparing concepts among individuals, concept similarity depends most on the similarity of core conceptual features than on less essential relations. There is not a hard distinction between the kinds of properties which can play a core role and those which play an inessential role; it is possible that the same properties which form the core of one person's private concept play only a small role in someone else's private concept.

Two examples will help illustrate this point. First, consider two students who are asked to recreate the proof of a theorem. Both students are able to perform the task successfully. However, they accomplish the task in two different ways. Student A has memorized the steps of the proof, and on the exam simply writes the steps from memory. Student B does not explicitly remember the steps of the proof, but understands the mathematical concepts necessary to complete the proof. It is obvious that the two students will differ in their cognitive phenomenologies, as they are performing entirely different conscious cognitive processes to achieve the same task. Student A's cognitive experience likely involves explicit memory recall of the steps of the proof, while Student B, in contrast, likely focuses primarily on the relevant mathematical concepts. Interestingly, this is true even if Student A has an understanding of the mathematical concepts, and Student B has explicitly learned the steps of the proof. As a result, some features forming the core of Student A's conscious experience are only on the fringe of Student B's, and vice versa. The major takeaway from this example is that the semantic holism underlying synchronic cognitive phenomenology does not depend on any sort of conceptual

essentialism; moreover, the network of semantic relations among beliefs should not be thought of as stable, but rather as highly context-dependent.

Second, consider two expert astronomers who are talking about Pluto. Though they share an enormous amount of expert knowledge about Pluto, Astronomer A is adamant that Pluto is not a planet, while Astronomer B insists that Pluto is in fact a planet. This case is interesting for at least two reasons. First, it must be apparent that generally speaking, there is much greater similarity between the concepts of Astronomer A and Astronomer B than between either astronomer and a layperson, regardless of what the layperson believes about Pluto's planetary status. The massive amount of expert knowledge shared by the two astronomers must outweigh the importance of the singular dispute of whether or not Pluto is a planet. Yet in the case where Astronomer A is explicitly debating the planetary status of Pluto with Astronomer B, this dispute becomes far more important. Suppose that the layperson agrees that Pluto is a planet. In the explicit debate about whether Pluto is a planet, it might plausibly be asserted that Astronomer B and the layperson share more similar concepts than the two astronomers with each other. These two points can only be reconciled if we assume that the core importance of certain related concepts may be context-sensitive. In general contexts, it does not make much of a difference if Pluto is a planet so long as the two astronomers agree on the massive amount of expert knowledge they share. Furthermore, because the layperson lacks this expert knowledge, their concept is radically different. However, in the explicit debate about Pluto's status, this point becomes more vividly in focus in each person's cognitive phenomenology. In that context it makes much more of a difference to the private concepts.

The contextual variability of cognitive phenomenal content is difficult to formalize. However, it seems that an individual may possess multiple different private concepts for the same object or property. In the above case, rather than thinking of the same private concept of Pluto changing radically from context to context, it is simpler to think of an individual employing different private concepts (or having different cognitive phenomenologies) for the same object under different aspects. Under the aspect of an object of scientific investigation, the astronomers share a sufficiently similar private concept, but under the aspect of a potential candidate for planetary status, their private concepts diverge significantly. Another useful example involves properties whose conditions of application vary by context. For example, the conditions for applying the concept ‘tall’ vary depending on whether one is discussing ordinary people, NBA players, or skyscrapers. This would be explained by the individual possessing different private concepts for ‘tall’ under different aspects: ordinary people, NBA players, skyscrapers, etc. Different people will likely disagree about exactly where the concept ‘tall’ applies even when discussing the concept under the same aspect. These different aspects are manifest in an individual’s dispositions to have particular cognitive phenomenologies in particular contextual conditions.

I discuss the criteria for comparing private concepts between individuals more directly in the next section. However, a deeper concern is already lurking. The Pluto example was innocuous because the relevant property typically does not make much of a difference to cognitive phenomenology when there is sufficient alternative similarity. Yet we can imagine cases involving massive theory change, in which two individuals may agree on a large number of empirical facts about some concept, but disagree about a property which makes a much more

fundamental difference to the concept. Skeptics may wonder if in such cases these private concepts can even be said to determinately correspond to the same public concept. I return to the concern of massive theory change in section VII.

V. Similarity of Cognitive Phenomenal Content

As discussed earlier, one of the greatest challenges for the theory of cognitive phenomenal holism is that it must avoid the concern of radical solipsism: it must be possible for individuals to communicate their thoughts to one another. In order for person A to successfully communicate some intentional content p to person B, person B must come to possess the same intentional content p' . This means that in order for communication to be possible, person A and person B must share the same intentional content. The goal of this section is to articulate what it means to share content in phenomenological terms.

Earlier, I discussed how many theories of mental content use propositions to avoid this concern. Propositions provide a neat account of communication: person A has a thought which expresses the proposition that p , person A utters a sentence in natural language which expresses the proposition that p , person B hears and interprets this utterance, and comes to have the thought which expresses the proposition that p . However, in that section I also demonstrated how that model necessarily contains the seeds of an argument for externalism, according to which the mental content will depend in part upon factors outside the individual, such as experts in the language community and the natural kinds which occur in the environment.

In contrast, I introduced the model of cognitive phenomenal holism, according to which one's mental content is determined entirely by the cognitive phenomenology occurring at a given

moment. One implication of this model is that mental content can vary widely between people who profess to be thinking and talking about the same content. In order for two people to literally share the same mental content, it must be the case that they have the same cognitive phenomenology; this would require that all their concepts were identically formed, that their cognitive phenomenal thresholds were the same, and that they were in the same relevancy context. The likelihood of this ever occurring seems impossible. As Pitt puts it:

“If cognitive intentionality is phenomenal, and phenomenal properties are not intersubjectively accessible, then knowledge of shared intentional content (and, hence, shared linguistic content), becomes (though not impossible) problematic. Though what you say you think provides me with evidence of what you do think, the possibility that the way I construe the evidence – the content I assign to your utterance – is not the content of your thought is a live one.”⁸¹

This implication of may sound problematic enough to compromise the entire theory. Yet as Pitt also says, “I don’t find any of this very disturbing.” To see why, let us first consider what it means to share perceptual content.

Imagine that you and I are both sitting on the beach. We feel the heat of the sun on our skin and the rough sand against our legs. We hear the rhythmic sound of the waves crashing onto the shore and the chirping of birds overhead. We smell the salt on the water, and we see the sparkle of sunlight against the deep blue of the ocean. Do we share all, or any, of the same perceptual phenomenology? In one sense, of course we do. We can discuss all of these perceptual features: I can say things like “Doesn’t the sun feel nice after a long winter?” and

⁸¹ Ibid.

“Isn’t this sand a little itchy?” If you agree, then it seems that we must share some perceptual phenomenal content. Yet in two important senses, we do not share the same content. First, our sensory inputs are not identical. You are sitting slightly to my left, and as a result, a slightly different configuration of sand presses against your legs. The angle from which you can see the ocean is different, and as a result, the sun glints at different times and locations for you. This slight difference in position results in a different phenomenology, though many of the general features are shared.

Yet even if we suppose that you and I are somehow in exactly the same position, such that our sensory inputs are identical, we still do not share the same content. Your experience at the beach depends in part on your internal phenomenological composition at the time. For example, if you have a headache, then the slight glint of light off the ocean may be deeply unpleasant, whereas I find it beautiful. If you have sensitive skin, the sand may feel coarse and abrasive, whereas I find it soft and soothing. You may be so distracted by the discomfort of the sand that you may fail to be aware of the crash of the waves altogether, while in my case, the rhythmic roar is all I can focus on. Furthermore, my past experiences at the beach will likely contribute to my phenomenology, and unless your past experience is identical to my own, the effect of this past experience will vary between us. Perceptual phenomenology is a function of one’s sensory inputs, one’s internal phenomenological state, and one’s past experience; it is unlikely that any two people ever share exactly the same phenomenal content.

One could take this to be a conclusion in favor of radical solipsism. We never have any way to compare our own experience to that of anyone else. Yet of course, we are able to compare our experiences. We have plenty of publicly ascertainable features by which we make

comparisons: the sun is warm, the sand is itchy, the horseradish is bitter, the stop sign is red. Of course, it is ultimately unknowable if anyone else is experiencing anything like what I am experiencing when I invoke these features (let alone that anyone else is experiencing anything at all). Still, because I know that other humans have a nervous system configured much like my own, it stands to reason that their experience is at least *sufficiently similar* to my own. This sufficient similarity is all I need to allow for communication of phenomenological experience.

The situation is even more optimistic for cognitive phenomenology. Just as in the previous case, the problem of other minds prevents me from knowing that anyone else shares my cognitive phenomenology (or that anyone else has cognitive phenomenology at all). Still, there are publicly ascertainable features according to which we can compare our thoughts, namely language. Recall Mendelovici and Bourget's example of the philosopher who uses the term 'supervenience.' When her colleague asks her to define the term, her colleague is making sure that their cognitive phenomenologies with regard to the term 'supervenience' are sufficiently similar so as to be about the same content. By giving a formal definition, both parties can make their cognitive phenomenologies with regard to the term more determinate, and bring each of their cognitive phenomenologies into alignment. In contrast, cases involving partially-grasped concepts, such as Bert's arthritis-concept, are cases in which mental content is insufficiently similar for communication to be possible. The principle of incompetence and cases of conceptual incoherence are grounds for insufficient similarity of cognitive phenomenal content.

VI. Massive Theory Change

A deeper problem for a radically internalist theory like cognitive phenomenal holism concerns massive theory change. In the Pluto case from earlier, the two astronomers had sufficiently similar concepts of Pluto so as to involve the same content, despite a significant conceptual difference. However, this sufficient similarity was largely dependent on a massive shared body of empirical knowledge about the planet.

Compare a modern meteorologist with an ancient Greek meteorologist. Likely, the ancient Greek meteorologist shares very little conceptual or empirical knowledge about lightning with the modern meteorologist. The concept of lightning is situated within a complex network of beliefs about the natural world, and the ancient and modern meteorologists' private concepts diverge at both the core and periphery of this network. One implication of this massive divergence is that they may not even be able to think of lightning under the same aspects. For example, the modern meteorologist can think of lightning as a movement of electrons. The ancient Greek meteorologist, lacking the prerequisite concept of electricity and subatomic particles, cannot think of the planet under this same aspect.

Still, there is a powerful intuition that the modern and ancient meteorologists are thinking about, or referring to, the same phenomenon. Their statements are made true or false by the same object. For this reason, externalist theories point to the shared causal origin of the concept (though Burge's theory is compatible with the view that the experts in each linguistic community, the astronomers themselves, define the concepts differently; this is in contrast with Putnam's natural kinds view and many causal-teleological views).

This problem is reminiscent of Loar's argument against functionalist internalist accounts of intentionality. Recall that Loar argued that intentionality is essentially outwardly directed; it at least "purports to refer" to external objects. He argues that functionalist accounts cannot provide a fully internalist account of demonstrative spatial concepts. In contrast, he claims that perceptual phenomenal experience provides the phenomenal basis for intentional content, which explains how intentional content "purports to refer."

Because my account is based on cognitive phenomenology, it denies that perceptual content is the basis for intentionality, as Loar suggests. However, I think Loar is on the right track. We can think of many concepts under the aspect of our basic perceptual experience. Both the modern and ancient meteorologist can think of lightning as the bright, linear pattern of light in the sky occurring during thunderstorms and typically preceding thunder. Crucially, the perceptual content itself is not the content of the cognitive phenomenology; rather, the perceptual content forms the basis for a core of conceptual content on which theoretical knowledge is based.

Prior to having any theory of thunder, one must pre-theoretically observe the phenomenon under the naive perceptual aspect described above; it is this pre-theoretical observation that calls into question the theory in the first place. Thus, the perceptually-based conceptual content at least temporally precedes the theory-laden concept. This provides the basis for something like Loar's suggestion of "intentional directedness," and this provides a universally accessible aspect for the possibility of sufficient similarity despite massive theory change. Some may object that once the theory is in place, even basic perceptual experience becomes theoretically laden; as a result, the modern and ancient astronomers can no longer access the shared pre-theoretical core once their theories are in place. However, I think this

objection misses the point: even if their actual perceptual experiences differ, they are entertaining the sufficiently similar content because they still are thinking about it under the same aspect. If their perceptual experiences were radically different (such as in the phenol-thio-urea case), then of course they would not provide a common core. I strongly doubt that theory-laden perception, if it is a real phenomenon, can change perceptual experience this radically; after all, one does not notice one's perceptual experience of lightning radically change after learning about electrons.

This phenomenon of intentional directedness helps tether private concepts to their external objects and provide stable truth conditions. This helps resolve a more general concern regarding truth conditions. One might think that private concepts theory makes large classes of false beliefs impossible. Consider Bert: under a Burgean interpretation, Bert falsely believes that he has arthritis in his thigh; under a private concepts interpretation, Bert truthfully believes that he has tharthritis in his thigh (and falsely believes that it is called 'arthritis'). The problem is that such interpretations rule out the possibility of an individual having a false conceptual belief, since every potential case can be reinterpreted as a true belief about a private concept. This problem may even extend to empirical beliefs. Suppose that Bert asserts that water boils at 32 degrees Fahrenheit. This could be interpreted as a false empirical belief about water, but it could also be interpreted as a true belief about a private concept of 'boiling,' where the concept aligns most closely to the phenomenon most people know as 'freezing'.

Because private concept theory already argued that this private concept interpretation is the right way to ascribe mental content in some cases, why not do so in every case? The theory needs a principled way to determine which cases require ascription of a private concept, and which involve false beliefs about a public concept. First, it is important to assert that there is a

fact of the matter to this dispute. It is not merely a matter of which interpretation we prefer; rather, the cognitive phenomenology determines which interpretation is correct. More precisely, the subject has a false belief about a public concept if her cognitive phenomenology is intentionally directed toward the object of the public concept; she has a true belief about the private concept if their cognitive phenomenology is intentionally directed toward something other than the object of the public concept.⁸² Bert has a true belief about tharthritis because the nature of his conceptual confusion makes his concept intentionally directed toward something other than the public concept of arthritis.

Another important point is that evaluation of truth conditions presupposes sufficient similarity of private concepts. At a minimum, it presupposes intentional directedness toward the same object. In the case involving ‘boiling,’ before we can determine what content to ascribe Bert, we must determine what he thinks about boiling. Suppose he asserts “Water boils at 32 degrees.” If someone asks, “What do you think boiling is?” and he responds “The transition of matter from a liquid state to a gaseous state,” then he likely has a false belief about a shared public concept; in contrast, if he asserts “the transition of matter from a liquid state to a solid state,” then he likely has a true belief about a private concept which (partially) aligns with the public concept ‘freezing’. As a result, we can inquire to better understand the cognitive phenomenology of our interlocutors, and evaluate truth with respect to cognitive phenomenologies brought into alignment.

⁸² This is not to give the public concept any ‘expert’ role in the language community; rather, the public concept is simply the thing collectively described by individuals’ private concepts intentionally directed toward the same object. I describe this in more detail in the next section.

VII. Public Concepts

Though this account of a private concept theory appears to avoid most of the major problems such theories might face, one might object that the reliance on public concepts to explain shared reference could lead to a reemergence of Burge's argument. Recall that the private concept theory explained the arthritis case in terms of whether or not Bert's private concept of arthritis was intentionally directed toward the same public concept as other members of his community. One might think that this 'public concept' begins to play the shared and objective role that propositions or expert knowledge played in the models of mental content targeted earlier. However, under a Burgean model of mental content, Bert's mental content could be *about* the public concept of arthritis even if Bert himself was conceptually confused. To Burge, because experts in the linguistic community are relied upon to know what arthritis is, personal concepts of arthritis refer to those public expert concepts.

Public concepts are not meant to work in such a way, nor must they in order to play the theoretical role they serve. First, public concepts are not defined by expert knowledge; rather, they consist in the general shared knowledge about a concept under an aspect. For example, in the case of arthritis, much of the knowledge an expert might have will not appear in the public concept (under the aspect of common parlance), as most competent non-expert users of the concept lack expert knowledge. As the Bert case reveals, one need not know *everything* about arthritis in order for one's mental content to be intentionally directed toward the public concept, but one must at least know that it is a condition of the joints. The public concept should be thought of as the minimally acceptable base concept for the users under the relevant aspect. Of

course, if a group of medical professionals were discussing arthritis, the minimal acceptable base concept under the aspect would include far greater expert knowledge.

Furthermore, public concepts may not always be particularly public. Public concepts only exist among groups of people whose private concepts align sufficiently to allow for minimal acceptable base concepts. However, when comparing the concepts of people from disparate groups, there may be no minimal acceptable base concept. It should be unsurprising that there may be concepts which are deeply embedded in certain cultures or groups, and which outsiders to the group may fail to understand. These should be understood as cases in which the two parties cannot find a minimally acceptable base concept as a starting point.

Imagine that your friend introduces you to a card game you have never heard of. As your friend begins to explain the mechanics of the game, some of it sounds very similar to a game with which you are already familiar. At this point, you must figure out whether your friend is describing the same game by a different name, or whether he is genuinely describing a new game. After hearing a sufficient number of rules, your understanding of the game becomes determinate enough to decide whether or not the game is new to you. Suppose it is in fact the same game under a different name. Your friend may possess a great deal more expert knowledge about the game, including perhaps the history of the game and nuanced strategy. Still, this does not prevent you and your friend from playing together. Finally, there may be certain rules which are only enforced in formal settings, such as tournament play, and not during casual games.

Public concepts are like card games. They do not exist independently of the individuals who know the rules and play the game, just as the public concepts do not exist independently of individuals with private concepts who aim to communicate with one another. Furthermore, there

is a minimal level of shared understanding necessary in order to be playing the game at all. Similarly, there is a minimally acceptable base concept in order for two parties' respective mental content to be intentionally directed toward the same public concept. Beyond this, further expert knowledge by one party does not prevent them from sharing a public concept, though it does change the cognitive phenomenal experience of the respective private concepts. Finally, what standards are minimally acceptable vary according to the aspect, just as certain elements of the game may depend on the degree of formality of the setting. The situation described above may occur in the case of concepts as well: it may be unclear whether I and my interlocutor are thinking of the same concept under the same name, or whether we are thinking of distinct concepts entirely. We can only resolve this indeterminacy by articulating our respective concepts more clearly. The degree of determinacy necessary to resolve the indeterminacy will again vary according to the aspect.

Public concepts are useful as a theoretical entity for two reasons: private concepts are radically variable, and language has a limited capacity to communicate mental content. We know that our mental lives are rich and determinate, and that each of our concepts is formed by a massive array of belief content, personal experience, and logical entailments; yet at the same time, we know that we are somehow able to communicate with one another. Public concepts are the resource which makes this possible. They are only precisely as determinate as they need to be; only as determinate as necessary to differentiate themselves from related concepts. Indeed, the reason that more formal or professional aspects require a greater degree of determinacy for a minimally acceptable shared base concept is that such settings require greater precision in differentiating between closely related concepts. Public concepts are only as good as they need to

be, and we must accept that a rich and massive amount of cognitive content remains unshared between individuals.

Bibliography

- Block, Ned. "An Argument for Holism." *Proceedings of the Aristotelian Society* 95 (1995): 151-164.
- Burge, Tyler. "Individualism and the Mental." *Midwest Studies in Philosophy* 4, no. 1 (1979): 73-122.
- Brandom, Robert. *Making it Explicit*. Cambridge: Harvard University Press, 1994.
- Chalmers, David. "The Representational Character of Experience." In *The Future For Philosophy*, edited by Brian Leiter. Oxford: Oxford University Press, 2004.
- Dretske, Fred. *Knowledge and the Flow of Information*. Cambridge: MIT Press, 1981.
- Naturalizing the Mind*. Oxford: Oxford University Press, 1995.
- "Phenomenal Externalism, or If Meanings Ain't in the Head, Where Are Qualia?" *Philosophical Issues* 7 (1996): 143-158.
- Horgan, Terence and John Tienson. "The Intentionality of Phenomenology and the Phenomenology of Intentionality." In *Philosophy of Mind: Contemporary and Classical Readings*, edited by David J. Chalmers. Oxford: Oxford University Press, 2002.
- Farkas, Katalin. "Phenomenal Intentionality Without Compromise." *The Monist* 91 no. 2 (2008), 273-293.
- Fodor, Jerry. "Semantics, Wisconsin Style." *Synthese* 59, no. 3 (1984): 231-250.
- Jackson, Frank. "Epiphenomenal Qualia." *Philosophical Quarterly* 32 (1982): 127-136.
- Kriegel, Uriah. "The Phenomenal Intentionality Research Program." In *Phenomenal Intentionality*, edited by Uriah Kriegel. Oxford: Oxford University Press, 2013.
- Loar, Brian. "Phenomenal Intentionality as the Basis of Mental Content." In *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, edited by Martin Hahn and B. Ramberg. Cambridge: MIT Press, 2003.
- Lycan, William. "Phenomenal Intentionalities." *American Philosophical Quarterly* 45, no. 3 (2008): 233-252.
- Consciousness and Experience*. Cambridge: MIT Press, 1996.
- "The Case for Phenomenal Externalism." *Philosophical Perspectives* 15 (2001): 17-35.

- Mendelovici, Angela and David Bourget. "Tracking Representationalism." In *Philosophy of Mind: The Key Thinkers*, edited by Andrew Bailey, 209-235. New York: Continuum, 2014.
- "Phenomenal Intentionality." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Stanford: The Metaphysics Research Lab, 2017.
- "Consciousness and Intentionality." In *Oxford Handbook of Consciousness*, edited by Uriah Kriegel. New York: Oxford University Press, forthcoming.
- Millikan, Ruth. "Biosemantics." *Journal of Philosophy* 86 (2009): 281-297.
- Nagel, Thomas. "What Is It Like to Be a Bat?" *The Philosophical Review* 83, no. 4 (1974): 438.
- Pitt, David. "The Phenomenology of Cognition, or What Is It Like to Think that P?" *Philosophy and Phenomenological Research* 1 (2004): 1-36.
- "Intentional Psychologism." *Philosophical Studies* 146, no. 1 (2009): 117-138.
- "Conscious Belief," *Rivista Internazionale di Filosofia e Psicologia* 7 (2009): 124.
- "Introspection, Phenomenality, and the Availability of Intentional Content." In *Cognitive Phenomenology*, edited by Tim Bayne and Michelle Montague. Oxford: Oxford University Press, 2011.
- "The Burgean Intuitions." Unpublished.
- Quine, W.V.O. "Two Dogmas of Empiricism." In *From a Logical Point of View*. Cambridge: Harvard University Press, 1951.
- Searle, John R. "Minds, Brains, and Programs." *The Behavioral and Brain Sciences* 3, no. 3 (1980): 417-424.
- "Consciousness, Unconsciousness and Intentionality." *Philosophical Issues* 1 (1991): 45-66.
- The Rediscovery of the Mind*. Cambridge: MIT Press, 1992.
- Stampe, Dennis. "Toward a Causal Theory of Linguistic Representation." *Midwest Studies in Philosophy* 2, no 1. (1977): 42-63.
- Tye, Michael. *Ten Problems of Consciousness*. Cambridge: MIT Press, 1995.
- Yli-Vakkuri, Juhani and John Hawthorne, *Narrow Content*. Oxford: Oxford University Press, 2018.