
Knowing What It's Like

Andrew Y. Lee | andrewyuanlee@gmail.com | <http://andrewyuanlee.com/>
University of Oslo, Department of Philosophy

WORD COUNT (including footnotes & references): 9088

KEYWORDS: phenomenal knowledge, phenomenal concepts, concept mastery, inexact knowledge, experience requirement, pure phenomenal concepts, revelation

ABSTRACT:

This paper develops a degree-theoretic account of knowledge of what it is like to have an experience. I argue (1) that knowledge of phenomenal character varies along a spectrum from the more exact to the more approximate, and (2) that phenomenal concepts vary along a spectrum in how precisely they characterize what it is like to undergo their target experiences. I motivate this degreed picture by appeal to limits in epistemic abilities such as recognition, imagination, and inference. I argue that approximate knowledge of phenomenal character cannot be explained merely by appeal to determinable or vague phenomenal concepts. I discuss how phenomenal concepts that yield more exact knowledge of what it is like to undergo their target experience are those that eliminate more “phenomenal possibilities.” And I explain how the resulting view challenges some common assumptions about the acquisition conditions, requirements for mastery, and referential mechanisms of phenomenal concepts.

Introduction

Consider what it is like to feel pain, to see red, or to smell cinnamon. Then consider what it is like to undergo the echolocation experiences of bats, the proprioceptive experiences of octopuses, or the electromagnetic experiences of aliens. There is an obvious asymmetry between your ability to think about the former versus your ability to think about the latter. What explains the asymmetry?

The standard explanation is that the difference is a matter of whether or not you possess phenomenal concepts for the relevant experiences. You know what it is like to feel pain, see red, and smell cinnamon because you have phenomenal concepts for those experiences, but you do not know what it is like to echolocate, to move your seventh tentacle spirally, or to sense a polarized magnetic field because you lack phenomenal concepts for those experiences.

This paper argues that the aforementioned asymmetry is a difference in degree, rather than a difference in kind. I argue (1) that knowledge of what it is like to have an experience varies along a spectrum from the more approximate to the more exact, and (2) that phenomenal concepts vary along a spectrum in how precisely they characterize what it is like to undergo their target experiences. The goal of this paper is to develop, motivate, defend, and explore the consequences of this degreed picture of phenomenal concepts and phenomenal knowledge

§1 explains the degreed picture and contrasts it with the all-or-nothing picture. §2 motivates the degreed picture by appeal to limits in epistemic abilities such as recognition, imagination, and inference. §3 argues that inexact knowledge of phenomenal character cannot be explained merely by appeal to determinable or vague phenomenal concepts. §4 develops a framework for systematizing the degreed picture, where inexactness of phenomenal knowledge is modeled by the number of possibilities ruled out for what it might be like to undergo the target experience. §5 explains how the degreed picture challenges some common assumptions about the referential mechanisms, acquisition conditions, and requirements for mastery of phenomenal concepts.

§1 | The Degreed Picture

A *phenomenal concept* is any concept of an experience that enables one to think about what it is like to have that experience (or, equivalently, the phenomenal character of that experience). If you think about what it is like to see red, feel pain,

or smell cinnamon, then you are deploying phenomenal concepts. Although every phenomenal concept is a concept of an experience, but not every concept of an experience is a phenomenal concept. This is because phenomenal concepts not only refer to experiences, but also enable one to think about what it is like to have those experiences. For example, the concept MY FAVORITE MENTAL STATE might happen to refer to an experience, but that does not make it a phenomenal concept.

The *target experience* of a phenomenal concept is the experience that the phenomenal concept refers to. Note that saying that x is the target experience of A is stronger than merely saying that x is in the extension of A : for example, scarlet is in the extension of RED, but RED is not a concept of scarlet (it is a concept of red). In general, I will use the term 'target experience' as neutral between phenomenal properties vs. particular experiences. To simplify the language, I will usually omit the term 'phenomenal' when talking about concepts, properties, and experiences.

I will take for granted that concepts are *mental representations*. More specifically, I assume that concepts are individuated by their psychological roles, are the constituents of thoughts, and enable epistemic abilities such as recognition, imagination, and inference. This view of concepts is dominant in cognitive science and is common amongst philosophers of phenomenal concepts.¹ On this framework, we can distinguish between concepts (i.e., mental representations that are constituents of thoughts), senses (i.e., the meanings of concepts), and referents (i.e., the extensions of concepts).² The main alternative framework construes concepts as *abstract entities*: in particular, as senses (rather than mental representations) that constitute propositions (rather than beliefs). Though I will speak in the language of the mental

¹ As examples, Chalmers [2003, p. 4] says "I take concepts to be mental entities on a par with beliefs: they are constituents of beliefs...analogous to the way in which words are constituents of sentences" and Balog [2012, p. 9] says that "concepts are mental representations that are...words of Mentalese."

² Some readers might think that I use the term 'concept' to mean what other philosophers sometimes mean by 'conception'. However, a *conception* is standardly defined as the set of beliefs associated with a concept. By contrast, I take concepts to be mental representations that are the constituents of thoughts, rather than sets of beliefs.

representation framework, those who prefer the abstract entity framework can translate my talk of concepts into talk of the mental representations used to grasp concepts.³

By possessing a phenomenal concept, one can acquire *knowledge of phenomenal character*, or knowledge of what it is like to undergo the target experience. My preferred view is that the locution 'knowing what it is like to have x ' (for some experience x) denotes a different kind of knowledge than locutions of the form 'knowing that P ' (for some phenomenal fact P) and 'knowing how to ϕ ' (for some experiential ability ϕ). I favor this view because I think that knowledge of phenomenal character is grounded in, rather than merely acquirable from, phenomenal concepts.⁴ However, the arguments of this paper are compatible (given some terminological substitutions) with views that instead take knowledge of phenomenal character to simply be a form of knowledge-that or knowledge-how, or that hold that there is merely a causal (rather than constitutive) relationship between phenomenal concepts and knowledge of phenomenal character. As a neutral locution, I will talk of phenomenal concepts *yielding* knowledge of phenomenal character, by which I mean possession of the phenomenal concept puts one in a position to acquire knowledge of what it is like to undergo the target experience.

According to the *degreed picture*, phenomenal concepts vary with respect to *degrees of purity*. The term 'degree of purity' denotes a theoretical role: what it is for a phenomenal concept to have a higher degree of purity is for it to yield more exact (as opposed to approximate) knowledge of what it is like to undergo its target experience.⁵ Equivalently, in my view, what it is for a phenomenal concept to have a

³ See Margolis & Laurence [2007] for discussion of these frameworks. There is also the view that concepts *are* epistemic abilities (rather than mental representations that enable epistemic abilities). The arguments of this paper apply straightforwardly to this view as well.

⁴ This is distinct from the idea that there is knowledge by acquaintance. In my view, mere acquaintance with an experience does not suffice for any kind of knowledge (since knowledge, but not experience, requires concept application), and one can have knowledge of phenomenal character even for experiences one has never had (a point I argue for in §3).

⁵ My term 'purity' stems from the now common distinction (introduced in Chalmers [2003]) between 'pure phenomenal concepts' (which refer to experiences directly via phenomenal character) and 'impure phenomenal concepts' (which refer via other means). However, under the degreed picture, (1) purity is a matter of degree, (2) phenomenal concepts with zero

higher degree of purity is for it to represent more precisely characterize what it is like to undergo its target experience. On the degreed picture, for any experience x , there are many phenomenal concepts of x that yield knowledge (at differing degrees of exactness) of what it is like to undergo x . If we grant the view I favored in the previous paragraph, then the question of whether the degreed picture is correct is equivalent to the question of whether there is inexact knowledge of phenomenal character.

By contrast, according to the *all-or-nothing picture*, for any experience x either you possess a phenomenal concept of x or not: if so, then you know what it is like to undergo x , and if not, then you do not. To get an initial grip on the difference between the degreed picture and the all-or-nothing picture, consider someone who knows what it is like to see red but does not know what it is like to see scarlet. Whereas the degreed theorist could say that this person has a phenomenal concept of scarlet that yields only approximate knowledge of what it is like to see scarlet, the all-or-nothing theorist must say that this person merely has a phenomenal concept of red (and perhaps factual knowledge that scarlet is a kind of red experience). At this point, the distinction may seem subtle, but a good deal of this paper will focus on elucidating the distinction and explaining its philosophical significance.

What evidence is there of philosophers favoring the all-or-nothing picture? A first source of evidence is that most discussions of phenomenal concepts merely distinguish phenomenal concepts from non-phenomenal concepts of experiences, with the implicit implication that there are no further important distinctions to be made within the class of phenomenal concepts and with no mention of the idea that phenomenal concepts exhibit the kind of degreed structure I have described. In fact, to my knowledge, there is no discussion of inexactness or degrees in any overview of either phenomenal concepts or phenomenal knowledge, and there is nearly no overlap between the literature on inexact knowledge and the literatures on phenomenal concepts and phenomenal knowledge.⁶

purity are distinct from non-phenomenal concepts of experiences (a point I argue for in §2), and (3) no phenomenal concepts refer directly via their phenomenal character (a point I argue for in §3).

⁶ A good deal of the phenomenal concepts/knowledge literature focuses on what makes phenomenal concepts/knowledge distinct from other kinds of concepts/knowledge and how

A second source of evidence is that philosophers sometimes explicitly state that phenomenal concepts yield exact knowledge of phenomenal character. For example, Chalmers [2003] says that when “Mary believes roses cause [red] experiences, or I am currently having [a red] experience, she thereby excludes all epistemic possibilities in which roses cause some other quality...or in which she is experiencing some other quality: only epistemic possibilities involving phenomenal redness remain.” In fact, this example is probative: the target experience in this instance is a determinable phenomenal property of phenomenal redness, yet the concept is nevertheless characterized as eliminating *all* epistemic possibilities for what its referent is like besides those involving red experiences. I will return to this idea in more detail later.

A third source of evidence is that philosophers working on phenomenal concepts/knowledge commonly endorse claims that we will later see are incompatible with the degreed picture. These include the claims (1) that phenomenal concepts refer to their target experiences directly via phenomenal character, (2) that phenomenal concepts can be acquired only by those who have had the relevant target experience, and (3) that phenomenal concepts enable one to know the essences of their target experiences.⁷ It may not yet be obvious why the degreed picture is in tension with these claims, but I will discuss these points in detail in §3.

that distinctness bears on the mind-body problem. For a limited sample of work on these issues, see Loar [1990], Sturgeon [1994], Hill [1997], Hill & McLaughlin [1998], Balog [1999], Perry [2001], Papineau [2002], Chalmers [2003], Levin [2006], Sundström [2011], and McLaughlin [2012]. For general overview, see Balog [2009], Nida-Rümelin & O’Connell [2019], and the papers in Alter & Walter [2006]. For discussions of inexact knowledge, see Williamson [1992], Mahtani [2008] and Carter [2019]. The philosophical literature that I think most directly connects to the issues in this paper is a small literature on phenomenal concepts and concept mastery, which I discuss in §3.

⁷ Other common commitments that are arguably in tension with the degreed picture, but which I do not discuss in detail in the paper, include the ideas that phenomenal concepts have modes of presentation that are identical to their referents (Loar [1990], Carruthers [2003], Tye [2003]), that phenomenal concepts are partially constituted by their target experiences (Papineau [2002], Balog [2012]), and that phenomenal concepts have identical primary and secondary intensions (Chalmers [2009], Goff [2011]).

There may be a temptation to think that many physicalists nowadays disavow the all-or-nothing picture. After all, most physicalists deny that phenomenal concepts enable one to know everything about their target experiences (since phenomenal concepts do not reveal the physical nature of their target experiences).⁸ However, the all-or-nothing picture concerns only what phenomenal concepts reveal about *what it is like to undergo* their target experiences. A physicalist may very well hold that phenomenal concepts do not reveal the physical nature of their target experiences yet still assume that phenomenal concepts yield exact knowledge of phenomenal character. The sources of evidence mentioned above indicate that many physicalists at least implicitly favor this kind of view.

Since few philosophers have explicitly discussed the idea of approximate knowledge of phenomenal character, some readers may wonder whether prior discussions of phenomenal concepts were merely idealizations intended to be compatible with the degreed picture (rather than implicit commitments in favor of the all-or-nothing picture). Though in some cases this may be plausible, I think the evidence cited above demonstrates that this hypothesis does not apply across the board. But even if we were to suppose that the degreed picture is what most philosophers have implicitly had in mind all along, it remains the case that the degreed picture has not been developed in detail and that its philosophical implications have been underappreciated. In light of this, this paper may likewise be interpreted as building upon prior work to develop a more nuanced picture of knowledge of what it is like.

§2 | A Study of Scarlet

To illustrate the intuitive appeal of the degreed picture, I will discuss three cases involving subjects whose phenomenal concepts arguably yield only approximate knowledge of phenomenal character. Each case concerns *scarlet experience*, which we can stipulate to be the kind of color experience normal humans have when looking at scarlet color chips under ideal conditions. Let us also stipulate that there is no more specific way of being a scarlet experience (so scarlet is maximally determinate) and that there are no borderline cases of scarlet experience (so SCARLET is not vague).

⁸ For some discussions of this form of physicalism, see Loar [1990], Balog [1999, 2012], Papineau [2002, 2006], Stoljar [2005], Chalmers [2007], and Pereboom [2011].

CASE 1: Ms. Scarlet has spent her life in a black and white room studying (but not having) color experiences. On day n , Ms. Scarlet's captors allow her to leave her room for five minutes to enter a new room. In this new room are one hundred color chips, each of which is a differing shade of red, each of which is labeled with the term for the kind of experience induced in Ms. Scarlet when she looks at that object, and one of which is scarlet (and labeled 'scarlet'). On each day after day n , Ms. Scarlet's captors allow her to reenter the new room for five minutes to look at the color chips. On each day, Ms. Scarlet also takes a test where she is asked to identify color experiences induced by unlabeled color chips. Before day n , her ability to recognize scarlet experiences is almost non-existent. On day $n+1$, her ability to recognize scarlet experiences is markedly better, though she still makes mistakes (such as categorizing a crimson experience as scarlet or categorizing a scarlet experience as vermilion). By day $n+100$, her ability to recognize scarlet experiences is extremely reliable, even when she is asked to identify scarlet experience against nearby red experiences.

If we follow conventional wisdom,⁹ Ms. Scarlet acquires a phenomenal concept of scarlet experience the very first time she leaves her room and sees the scarlet color chip. But what explains the changes in her epistemic abilities on the subsequent days? It is natural to think that on day n Ms. Scarlet knows only approximately what it is like to see scarlet, while by day $n+100$ Ms. Scarlet knows exactly what it is like to see scarlet. Since Ms. Scarlet's epistemic abilities gradually improve from day n to day $n+100$, and since she already possesses a phenomenal concept of scarlet experience by the end of day n , it follows that the epistemic changes cannot be explained merely by whether Ms. Scarlet possesses a phenomenal concept of scarlet experience. Instead, it seems that from day n to day $n+100$, Ms. Scarlet's phenomenal concept of scarlet experience changes so as to yield increasingly exact knowledge of what it is like to see scarlet.

CASE 2: Mr. Rainbow and Mr. Gray are professors who study color experience. Mr. Rainbow, moreover, has excellent epistemic abilities with respect to color experience: for example, he can imagine scarlet experience precisely and vividly, and he can acquire knowledge of many phenomenal facts about scarlet experience

⁹ I say 'conventional wisdom' in reference to standard views about Mary from Jackson [1982], whose situation parallels that of Ms. Scarlet (at least, until day n).

just by thinking about what it is like to see scarlet. Mr. Gray, on the other hand, has monochromacy (and so has never had a scarlet experience): nevertheless, he still sometimes tries to imagine what it is like to see scarlet (and imagines it as a kind of chromatic experience), and he can still know on that basis that what it is like to see scarlet is more similar to what it is like to see gray than what it is like to hear a trumpet. Though both Mr. Rainbow and Mr. Gray possess concepts of scarlet experience, Mr. Rainbow's concept arguably yields more exact knowledge of what it is like to see scarlet.

It may be tempting to argue that Mr. Gray simply lacks a phenomenal concept of scarlet experience. But consider Mr. Black, who is an equally competent expert on color experiences but who has never had any visual experiences (Mr. Black does not even have eyes). If neither Mr. Gray nor Mr. Black possesses a phenomenal concept of scarlet experience, then (by definition) neither is able to think about what it is like to see scarlet. However, Mr. Gray arguably has a better grasp than Mr. Black of what it is like to see scarlet (even though neither knows what it is like to see scarlet as well as Mr. Rainbow does). If Mr. Gray has some knowledge of what it is like to see scarlet, then Mr. Gray must be able to think about what it is like to see scarlet, from which it follows that Mr. Gray must possess a phenomenal concept of scarlet experience. Still, Mr. Gray's phenomenal concept of scarlet experience yields less exact knowledge of what it is like to see scarlet than Mr. Rainbow's phenomenal concept of scarlet experience, so Mr. Gray's phenomenal concept of scarlet experience is less pure than Mr. Rainbow's.

CASE 3: Consider your own knowledge of what it is like to see scarlet. You can reliably recognize instances of scarlet when presented against dissimilar experiences (such as non-red experiences), you can imagine what it is like to see scarlet (to at least some degree of precision and vivacity), you can know that seeing scarlet is similar to seeing other shades of red just by thinking about the phenomenal characters of those experiences, and you have a better grasp of what it is like to see scarlet than what it is like to echolocate. Therefore, you have at least some knowledge of what it is like to see scarlet. However, I suspect you cannot reliably recognize scarlet experience when presented against extremely similar red experiences, you cannot imagine exactly what it is like to experience scarlet (as opposed to other nearby red experiences), you cannot know simply on the basis of your phenomenal concepts that scarlet experience is exactly as similar with respect to

hue to crimson experience as it is to amaranth experience, and you are arguably in a better position to know what it is like to see scarlet if you are actually undergoing a scarlet experience than if you are merely thinking about scarlet experience. Therefore, you do not know exactly what it is like to see scarlet.

My appeal to these cases draws upon the assumption that limits in our recognitional, imaginative, and inferential abilities are evidence that our phenomenal concepts yield only approximate knowledge of what it is like to undergo their target experiences. Since we often cannot recognize target experiences with perfect reliability, imagine target experiences with perfect detail, or know all phenomenal facts about target experiences just on the basis of thinking about those target experiences, we have reason to think that our phenomenal concepts yield only approximate knowledge of phenomenal character. While it is possible to reject this connection between our epistemic abilities and our phenomenal concepts, doing so leaves one in an awkward position: if our phenomenal concepts enable us to know exactly what it is like to undergo their target experiences, then why do the associated epistemic abilities have a graded structure?

Though scarlet experience is my focal example, it is easy to see that these arguments generalize to phenomenal concepts for other experiences as well. For other examples, consider the maximally determinate total experience you had upon first waking up this morning, the complex flavor experience you have when drinking an interesting beer, or the visual experience you have when looking at a noisy mosaic of pixels. As before, your concepts of those experiences yield some knowledge of what it is like to have each of those experiences, given your recognitional, imagination, and inferential abilities. But as before, it is also plausible that you do not know exactly what it is like to have each of those experiences, given the limits in those epistemic abilities.

§3 | Objections

The cases above aimed to elicit the intuitive appeal of the degreed picture. Let me now address three common objections.

The Constitution Objection

The degreed picture seems in tension with the idea that thinking about an experience involves undergoing the experience that one is thinking about.¹⁰ For example, one might argue that when one deploys the concept RED, one undergoes a red experience. If one cannot even think about an experience without actually undergoing that experience, then how could it be that phenomenal concepts yield merely approximate knowledge of phenomenal character?

Now, few philosophers believe that one literally cannot think about what it is like to undergo an experience unless one is actually undergoing that exact experience. You can think about what it is like to be in severe pain without actually experiencing severe pain; if you are faced with the choice of either thinking about pain or experiencing pain, it is obvious which option is better. The more plausible idea is that thinking about an experience requires instantiating a mental image that resembles (but is phenomenally distinct from) the target experience. But that is compatible with the degreed picture, since there is no obvious reason for holding that mental images that merely resemble their target experiences yield exact knowledge of what it is like to have those target experiences.

It may be tempting to respond by appealing to direct phenomenal concepts, or concepts of occurrent experiences that are partially constituted by those occurrent experiences.¹¹ Those who accept the existence of direct phenomenal concepts might then argue that direct phenomenal concepts yield exact (rather than merely approximate) knowledge of phenomenal character. However, the degreed picture does not claim that *no* phenomenal concepts yield exact knowledge of their target experiences. The idea that phenomenal concepts vary in degrees of purity is compatible with thinking that there are special limit cases that have maximal purity. If direct phenomenal concepts exist, then they are candidates for those limit cases.

¹⁰ See Papineau [2002] and Balog [2012] for examples of views of this kind. Balog [2009] characterizes such views as involving “the idea that phenomenal concepts are constituted by the phenomenal experiences they refer to.”

¹¹ See Chalmers [2003] and Horgan & Kriegel [2007] for arguments in favor of direct phenomenal concepts. See Sundström [2011] for arguments against.

The Determinability Objection

Since scarlet is a determinate of red (and red a determinable of scarlet), it may be tempting to think that what I call a phenomenal concept that yields approximate knowledge of scarlet experience is really a phenomenal concept that yields exact knowledge of red experience. If this view is correct, then you do not possess a phenomenal concept of scarlet experience at all: instead, you possess merely a phenomenal concept of red experience.

As an initial response, note that we do not generally impose such demanding conditions on concept possession. Consider how one's concepts ARTHRITIS and ELM TREE and WEIGHT can refer to arthritis and elm tree and weight even if those concepts do not yield knowledge that arthritis is a disease of the joints, or that elm trees look the way they do, or that weight is an extrinsic property. The objector might counter that there are asymmetries between phenomenal concepts and other kinds of concepts. But even if we accept that there are such asymmetries (an issue that will be discussed in §3), the current objection still leads to counterintuitive consequences. If you do not possess a phenomenal concept of scarlet experience, then (by definition) you cannot think about what it is like to see scarlet. Yet on the face of it, you have been thinking about what it is like to see scarlet the entire time you have been reading this section. What else might you have been thinking about as you considered the cases discussed earlier? Instead of saying that you read paragraph after paragraph about scarlet experience yet were never thinking about scarlet experience, it is more natural to hold that you thought inexactly about what it is like to see scarlet as you read about scarlet experience.

Is there some countervailing reason for denying that you have been thinking about scarlet experience? A first response is that your phenomenal concept of scarlet experience has a mode of presentation with a coarse-grained content: for example, perhaps the mode of presentation represents the target experience only as some form of red experience. However, this response is more naturally understood as a specific theory about degrees of purity rather than as an objection to the degreed picture. A second response is that when you think that you are thinking about what it is like to see scarlet, you are really thinking about what it is like to see red and drawing an inference from your factual knowledge that scarlet is a type of red experience. But while that may be one way of acquiring knowledge that scarlet is a type of red experience, it is also plausible that you can simply think about what it is

like to see scarlet without drawing inferences from your beliefs about the relationship between scarlet experience and red experience. Unless we are systematically mistaken about the inferential structure of these conscious mental processes, this response overintellectualizes the psychological story.

The Vagueness Objection

A concept is *vague*¹² just in case it has borderline cases and *sharp* just in case it is not vague. For example, the concept BALD is vague, while the concept PHOTON is sharp.¹³ Since vague representations are inexact, it may be tempting to think that purity is just a matter of vagueness.

While there is a sense in which both vagueness and purity are a matter of inexactness, the nature of the inexactness differs: vagueness essentially involves borderline cases, whereas purity is independent of borderline cases. When I earlier defined the term 'scarlet experience', I stipulated that there are no borderline cases of scarlet experience: any color experience is either determinately a scarlet experience or not. If there are no borderline cases of scarlet experience, then any concept of scarlet experience must be sharp, since what it is for a concept to be vague is for that concept to admit of borderline cases. Yet even though concepts of scarlet experience must be sharp, they may nevertheless fail to be maximally pure. After all, I argued that your own concept of scarlet experience is not maximally pure and that other subjects (such as Mr. Rainbow and Mr. Gray) have phenomenal concepts of scarlet experience that vary in degree of purity. Since purity can vary even when vagueness is held fixed, purity is distinct from vagueness.

Is there vagueness without impurity? Suppose that persimmon experience is a borderline case of red experience, that you are as competent in thinking about red experience as one could possibly be, and that you know exactly what it is like to

¹² I assume here that vagueness is a semantic (rather than epistemic) phenomenon. The epistemicist about vagueness has analogous reasons for disentangling purity from vagueness, but it is more cumbersome to do so within an epistemicist framework.

¹³ Note that while vagueness is also commonly characterized as a property of *terms*, purity cannot be characterized as a property of terms. For example, there is no sense in which the term 'scarlet experience' is more or less pure than the term 'red experience'. This is further evidence that purity is distinct from vagueness.

see persimmon. Then your concept of red experience is vague since it has borderline cases yet also maximally pure since it yields maximal knowledge about what it is like to see red. As a contrast case, consider a colorblind person who has a phenomenal concept of red experience (for the kinds of reasons argued for earlier) but whose phenomenal concept does not even enable them to know that persimmon experience is a borderline case of red experience. The colorblind person's concept of red experience is as vague as (but less pure than) your concept of red experience. Therefore, there is a double dissociation between purity and vagueness.

§4 | Phenomenal Possibilities

In what follows, I develop a framework that provides a more systematic way of thinking about the knowledge of phenomenal character yielded by a given phenomenal concept. The core idea is that all phenomenal concepts rule out some (and leave open other) "phenomenal possibilities," and the more phenomenal possibilities ruled out by a phenomenal concept, the more exact the knowledge of phenomenal character yielded by that phenomenal concept.

Let a *phenomenal possibility* be a way that it might be like to undergo a target experience. Putting it another way, phenomenal possibilities are candidates for target experiences. Every phenomenal concept rules out some (and leaves open other) phenomenal possibilities. For example, your phenomenal concept SCARLET rules out the possibility that what it is like to experience scarlet is what it is (in fact) like to experience pain, but (given the arguments from earlier) it does not rule out the possibility that what it is like to experience scarlet is what it is (in fact) like to experience crimson. Now, there is an interesting question of how to best analyze the notion of a phenomenal possibility. Just as we might ask which kind of entity best satisfies the theoretical roles associated with the notion of a possible world, we might ask which kind of entity best satisfies the theoretical associated with the notion of a phenomenal possibility.¹⁴ For our purposes, though, it will be simplest to just think of

¹⁴ I suspect that the best analysis characterizes phenomenal possibilities as fuzzy sets of possible experiences, where a *possible experience* is a maximally specific way that a total experience could be. Under this approach, determinability would be modeled by the number of possible experiences contained within a phenomenal possibility, and borderline cases would

phenomenal possibilities as phenomenal properties. More specifically, we can think of the set of phenomenal possibilities as equivalent to the set of all phenomenal properties. A key advantage of this approach is that it provides a straightforward way of capturing determinability and vagueness (alongside purity). For example, the fact that SCARLET is more determinate than RED is captured by the fact that SCARLET leaves open more determinate phenomenal possibilities than RED. Likewise, the fact that RED is vague is captured by the fact that RED leaves open some phenomenal possibilities with borderline cases, such as persimmon experience.

Under this framework, the degree of purity of a phenomenal concept can be specified as the proportion of the set of all phenomenal possibilities ruled out by that phenomenal concept.¹⁵ This enables us to assign every phenomenal concept a *purity value* between 0 and 1 (corresponding to the proportion of phenomenal possibilities ruled out), where higher values indicate higher degrees of purity.¹⁶ Over the rest of the paper, I will denote phenomenal concepts by designating their target experiences in small-caps and their purity value in subscripts. For example, SCARLET_{0.2} denotes a relatively impure phenomenal concept of scarlet while RED₁ denotes a maximally pure phenomenal concept of red.

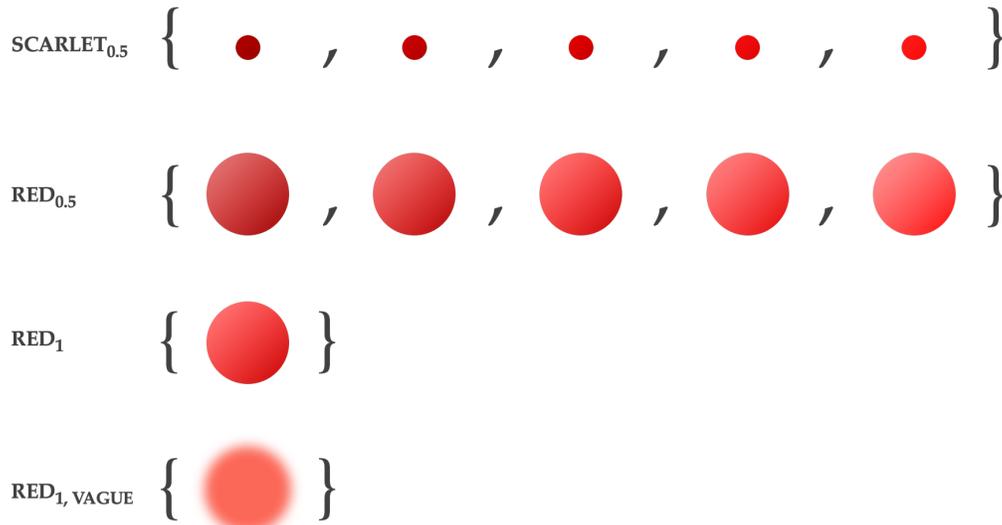
Now we can specify the limit cases for purity. At one limit are *maximally pure* phenomenal concepts, which have purity value 1, which rule out all phenomenal possibilities except one, and which yield exact knowledge of what it is like to have the target experience. At the other limit are *minimally pure* phenomenal concepts, which have purity value 0, which rule out no phenomenal possibilities (but still specify that the target experience is an experience), and which yield maximally approximate knowledge of what it is like to have the target experience. Between the extremes are *partially pure* phenomenal concepts, which have purity values between 0 and 1, which rule out some (but not all) phenomenal possibilities, and which yield approximate knowledge of what it is like to undergo the target experience.

be modeled by the degree of set membership for possible experiences. However, for limits of space (and to maintain focus on purity), I will not discuss this approach in detail.

¹⁵ See Tao [2011] for discussion of measures on infinite sets.

¹⁶ For notational convenience, assume that we scale purity values logarithmically, so that a purity value of $\frac{1}{2}$ denotes a moderately pure phenomenal concept (rather than an extremely impure phenomenal concept). Otherwise, a purity value of $\frac{1}{2}$ would mean leaving open half of all phenomenal possibilities, which presumably ought to count as extremely impure.

The framework also enables us to more systematically disentangle purity, determinability, and vagueness. Whereas purity is a matter of the number of phenomenal possibilities ruled out, determinability is matter of how specific those phenomenal possibilities are, and vagueness is a matter of the extent to which those phenomenal possibilities have borderline cases. These structural differences between purity, determinability, and vagueness are illustrated in the diagram below: each row contains the name of a phenomenal concept and an illustration of the set of phenomenal possibilities left open by that phenomenal concept, and each circle represents a phenomenal possibility, where the size of the circle represents degree of determinability and fuzziness represents borderline cases:



Purity vs. determinability vs. vagueness.

In brief: SCARLET_{0.5} and RED_{0.5} differ in determinability (but not purity or vagueness), RED_{0.5} and RED₁ differ in purity (but not determinability or vagueness), and RED₁-sharp and RED₁-vague differ in vagueness (but not purity or determinability). The most important contrast case, for our purposes, is RED_{0.5} versus RED₁. Both concepts are equally determinate since they leave open equally specific phenomenal possibilities. Both concepts are equally vague since neither has phenomenal

possibilities with any borderline cases. And yet RED_1 leaves open only one phenomenal possibility while $RED_{0.5}$ leaves open five, meaning RED_1 is purer than $RED_{0.5}$.

Finally, the framework enables us to specify the epistemic relations between different phenomenal concepts. If phenomenal concepts A and B leave open exactly the same phenomenal possibilities, then they yield the exact same knowledge of the phenomenal character of their target experiences. If A and B leave open overlapping sets of phenomenal possibilities, then for all one knows they might pick out target experiences that are phenomenally identical. If A and B leave open completely disjoint sets of phenomenal possibilities, then anyone who possesses A and B is in a position to know that the target experience of A is phenomenally distinct from the target experience of B. Last, if set of phenomenal possibilities left open by A is a proper subset of the set of phenomenal possibilities left open by B, then A eliminates all phenomenal possibilities that B eliminates plus more. In such a case, A is *strictly purer* than B, meaning that A yields strictly more exact knowledge than B of what it is like to undergo the target experience.

§5 | Philosophical Implications

The rest of this paper discusses implications of the degreed picture for questions concerning the referential mechanisms, acquisition conditions, and requirements for mastery of phenomenal concepts. Since space is limited, each of these discussions will be brief: the main purpose of this section is to illustrate some of the ways in which the degreed picture is philosophically significant.

Concept Acquisition

I began this paper by contrasting our knowledge of feeling pain, seeing red, and smelling cinnamon with our knowledge of the echolocation experiences of bats, the proprioceptive experiences of octopuses, and the electromagnetic experiences of aliens. The standard explanation of this asymmetry is that we possess phenomenal concepts of the former (but not the latter) experiences. A central aim of this paper has been to argue that that standard explanation oversimplifies.

On the degreed picture, we have relatively exact knowledge of what it is like to feel pain, to see red, and to smell cinnamon because we have relatively pure phenomenal concepts of those experiences. By contrast, we have extremely approximate knowledge of what it is like to echolocate, to move one's seventh tentacle spirally,

and to sense a polarized magnetic field because we have extremely impure phenomenal concepts of those experiences. The asymmetry is not a matter of whether we possess phenomenal concepts for the relevant experiences, but instead a matter of how pure our phenomenal concepts are.

Some readers might object that it is counterintuitive to claim that we possess phenomenal concepts even of bat, octopus, and alien experiences. Now, I suspect that most of the counterintuitive force is residue from the all-or-nothing picture. But setting that aside, the core aim of this paper is to elucidate the epistemic structure of phenomenal concepts (rather than to dictate how we apply the label 'phenomenal concept'). Suppose we reserve the term 'phenomenal concept' for concepts of experiences that surpass a certain threshold of purity, yielding the result that we lack phenomenal concepts for bat, octopus, and alien experiences. Nevertheless, it remains the case that the difference between our knowledge of the familiar experiences of normal humans versus our knowledge of the exotic experiences of other kinds of creatures is a matter of degree.

Are there any experiences for which we simply cannot acquire a phenomenal concept? A phenomenal concept is, by definition, a concept that enables one to think about what it is like to have an experience. The least pure phenomenal concepts yield knowledge only that there is something it is like to have the experience (with no further specificity on the particular phenomenal character).¹⁷ These minimally pure phenomenal concepts eliminate no phenomenal possibilities (but still characterize the target experience as an experience). Since we can represent any experience whatsoever as being such that there is something it is like to be in it, we can always acquire at least a minimally pure phenomenal concept of any experience.

This may raise the worry that the degreed picture makes the acquisition conditions for phenomenal concepts trivial. But notice that there remains a significant difference between concepts that represent experiences *as* experiences (i.e.

¹⁷ Why not instead require that phenomenal concepts eliminate at least one phenomenal possibility? Let A be a concept that eliminates one phenomenal possibility, B be a concept that eliminates zero phenomenal possibilities, and C be a concept that does not even represent its referent as an experience. Intuitively, the difference between A and B is much less significant than the difference between B and C. Because of this, I believe my definition of minimally pure phenomenal concepts better carves at the conceptual joints.

phenomenal concepts) versus concepts that refer to experiences but do not represent them as experiences (i.e. non-phenomenal concepts of experiences). Though zombies can acquire non-phenomenal concepts of experiences (since it is relatively easy to acquire a concept whose referent happens to be an experience), zombies cannot acquire even minimally pure phenomenal concepts (at least if we grant the plausible assumption that zombies cannot think about what it is like to have an experience). On the degreed picture, there is a smooth transition from maximally pure phenomenal concepts to minimally pure phenomenal concepts, at which point we cross the threshold to non-phenomenal concepts of experience. This mirrors a smooth transition from maximally exact knowledge to maximally approximate knowledge, at which point we cross the threshold to no knowledge of phenomenal character at all.

A noteworthy consequence concerns the idea that there is an experience requirement on the acquisition of phenomenal concepts: that in order to acquire a phenomenal concept of an experience, one must have undergone that experience (or another experience that is relevantly similar).¹⁸ Against this, I have argued that it is possible to acquire phenomenal concepts for all sorts of experiences one has never had. Though the experience requirement appears plausible if we presume the all-or-nothing picture, the requirement is less compelling once we adopt the degreed picture. The grain of truth in the experience requirement is that one can acquire much purer phenomenal concepts for experiences one has had than for experiences one has never had. But since possessing a phenomenal concept for an experience does not require knowing exactly what it is like to have that experience, the experience requirement is false.

Concept Mastery

This paper has assumed that concepts are mental representations. By contrast, discussions of concept mastery standardly take concepts to be abstract

¹⁸ See, for example, Sturgeon [1994], Tye [1995], and Papineau [2006] for endorsements of the experience requirement. See Ball [2009] for arguments against the experience requirement. Note that Ball builds the experience requirement into the definition of 'phenomenal concept', and so concludes that there are no phenomenal concepts. By contrast, I think it is more theoretically useful and reflective of actual use to define phenomenal concepts as concepts of experiences that enable one to think about what it is like to have those experiences.

entities.¹⁹ Under the abstract entity framework, many different mental representations can be used to grasp the same concept, which yields a natural distinction between merely possessing a concept versus mastering a concept. Under the mental representation framework, however, a change in one's mental representation often means a change in the concept itself. This makes it somewhat awkward to even formulate questions about concept mastery within the mental representations framework, and makes it easy for those of us who favor the mental representation framework to overlook issues about concept mastery.

Under the degreed picture, it is natural to think that concept mastery requires possessing a maximally pure phenomenal concept. If your phenomenal concept of scarlet experience does not enable you to know exactly what it is like to see scarlet, then it is plausible that you have not yet mastered thinking about scarlet experience. This hypothesis aligns with more general accounts of concept mastery, which often analyze mastery in terms of the endorsement or recognition of certain beliefs or inferences: if one's phenomenal concept does not yield exact knowledge of what it is like to have a target experience, then it is plausible that one's phenomenal concept will not permit (or enable) one to endorse (or recognize) the beliefs or inferences diagnostic of mastery.²⁰ Conversely, it is hard to see where to draw the line if one were to reject maximal purity as a requirement for mastery. After all, it is plausible that at least some knowledge of what it is like to have an experience is necessary for mastery in thinking about that experience. But once we accept that some purity is needed, there seems no principled cutoff short of maximal purity.²¹

The notion of concept mastery is useful for clarifying the explanatory ambitions of this paper. The degreed picture can be understood as providing an account of how phenomenal concepts yield knowledge of phenomenal character at all levels of mastery. In fact, I suspect that part of the reason the degreed picture has been

¹⁹ See Burge [1979] and Peacocke [1992] for some classic discussions addressing concept mastery. See Rabin [forthcoming] for a recent discussion. For discussions of mastery for phenomenal concepts, see Ball [2009, 2013], Rabin [2011], and Alter [2013]. Note that these discussions mostly focus on the ramifications of concept mastery for the knowledge argument, rather than on the kinds of issues addressed in this paper.

²⁰ See Rabin [forthcoming] on general accounts of concept mastery.

²¹ Note that I am suggesting possession of a maximally pure phenomenal concept only as a necessary (rather than sufficient) condition for mastery.

neglected is because philosophers working on phenomenal concepts tend to tacitly focus on idealized subjects that have achieved concept mastery. This hypothesis is supported by the fact that the philosophical literature on phenomenal concepts tends to adopt the mental representations framework of concepts (which tends to obscure issues about concept mastery, as noted above).

A noteworthy consequence concerns *revelation*, the thesis that possession of a phenomenal concept enables one to know the essence of the target experience.²² Suppose we grant that the essence of an experience is simply the phenomenal character of that experience. Then it may be tempting to think that revelation is true, especially if one considers only idealized subjects that have mastery over all their concepts. However, once we adopt the degreed picture (and broaden the scope to subjects who have not yet achieved mastery in thinking about experiences), revelation becomes less appealing. Just because one possesses a phenomenal concept of an experience x does not mean that one knows exactly what it is like to undergo x . And if one does not know exactly what it is like to undergo x , then one does not know the essence of x . Therefore, while revelation might be true for the special class of maximally pure phenomenal concepts, it is false for all other phenomenal concepts (and false simpliciter).

Reference

Let the *phenomenal content* of a phenomenal concept be the way that the phenomenal concept characterizes what it is like to undergo its target experience.²³ A popular idea in the phenomenal concepts literature is that the phenomenal content of a phenomenal concept suffices for determining reference to a target experience. This idea is more often expressed as the claim that phenomenal concepts refer to target experiences by directly specifying the target experiences' phenomenal

²² See Goff [2015] and Broi [2020] for some recent discussions of revelation. Note that while discussions of revelation are usually focused on the significance of the thesis for physicalism, my points here are intended to be neutral on questions about physicalism.

²³ I will take for granted that we can coherently talk of the contents of concepts. On this way of thinking, contents need not be propositional (since concepts cannot be true or false) and contents are distinct from referents (since concepts have content even when they fail to refer). See Weiskopf [2009] and Margolis & Laurence [2014] for more on the contents of concepts.

characters, or as the claim that phenomenal concepts have non-contingent modes of presentation that designate the same referents across internal duplicates.

The degreed picture is in tension with the aforementioned claims about reference. In order for that picture of reference to work, it must be the case that phenomenal concepts specify exactly what it is like to undergo their target experiences. Otherwise, the phenomenal content of a phenomenal concept would underdetermine which target experience is its referent.²⁴ Since the degreed picture entails that any phenomenal concept that is not maximally pure leaves open multiple phenomenal possibilities, and since phenomenal possibilities are candidates for what it is like to undergo target experiences, it follows that non-maximally pure phenomenal concepts leave open multiple candidates what it is like to undergo their target experiences. By consequence, the degreed picture entails that phenomenal contents those do not generally suffice for reference to particular target experiences.

Is this plausible? Consider again what it is like to see scarlet, and then consider what it is like to see vermillion. Speaking for myself, it is not clear that there is any difference between how my concept SCARLET characterizes what it is like to see scarlet versus how my concept VERMILLION characterizes what it is like to see vermillion. If we were to test my recognitional, imaginative, and inferential abilities, it is not clear that such tests would reveal any difference in how I think about what it is like to see scarlet versus what it is like to see vermillion.²⁵ Yet the arguments from §1 indicate that there is good reason to think that I can still think about what it is like to have each experience. So, while it is obvious that SCARLET refers to scarlet

²⁴ As examples, Loar [1990] says phenomenal concepts have “non-contingent phenomenal modes of presentation,” Chalmers [2003] says a phenomenal concept picks out its referent “directly, in terms of its intrinsic phenomenal nature,” Tye [2003] says “phenomenal concepts refer directly” in that they “have no associated reference-fixers, no descriptive content at all,” and Papineau [2007, p. 104] says “phenomenal concepts refer to phenomenal properties directly, and not by invoking any further contingent properties of those referents.”

²⁵ This does not mean the concepts are identical. My claims concern only phenomenal contents, or how the concepts represent *what it is like* to undergo the target experiences. This is compatible with thinking that the concepts differ in other respects: for example, perhaps my concept SCARLET (versus VERMILLION) enables me to know that scarlet (rather than vermillion) experience is normally caused by scarlet (as opposed to vermillion) color chips.

experience and VERMILLION refers to vermillion experience, it is unobvious whether the phenomenal contents of SCARLET and VERMILLION differ at all.

In fact, even the phenomenal contents of maximally pure phenomenal concepts arguably do not wholly determine reference. Recall that phenomenal concepts can refer to either phenomenal properties or to particular experiences. Suppose one has a phenomenal concept that is maximally pure, so it enables one to know exactly what it is like to undergo its target experience. Nevertheless, there remains a question of whether the phenomenal concept refers to a particular experience or to a maximally determinate total phenomenal property. And even if we were to set aside phenomenal properties and focus only on particular experiences, it is possible for there to be distinct particular experiences that are phenomenally identical. These observations indicate that the phenomenal contents of phenomenal concepts invariably underdetermine reference to their target experiences.

If phenomenal contents do not determine reference for phenomenal concepts, then how do phenomenal concepts refer? The natural move is to hold that phenomenal concepts refer via the same kinds of referential mechanisms as other concepts: for example, by definite descriptions, speaker intentions, causal chains, demonstrative applications, deference to experts, rules of use, or some combination of these mechanisms. Since these ideas are familiar, and since there is no obvious reason to expect anything idiosyncratic in how these referential mechanisms work for phenomenal concepts, I will not evaluate which of these accounts is most plausible.²⁶ The more interesting question is whether an appeal to these general referential mechanisms is sufficient for explaining how phenomenal concepts refer.

Although the degreed theorist must abandon the audacious claim that phenomenal contents wholly determine reference to target experiences, they can still satisfy the intuition that phenomenal concepts have special referential mechanisms. In particular, a degreed theorist might think that phenomenal contents constrain (but do not fully determine) reference, meaning that the referent of a phenomenal

²⁶ See Michaelson & Reimer [2019] for general discussion of theories of reference. See Ball [2009] and Rabin [manuscript] for more detailed discussion of how these kinds of kinds of referential mechanisms can be applied to phenomenal concepts. Note that while Ball's stated aim is to argue against the existence of phenomenal concepts, this difference between his view and my view may be verbal (see footnote 23).

concept must be amongst the phenomenal possibilities the concept leaves open. On this view, phenomenal contents cull the initial candidates for target experiences, while the referential mechanisms mentioned above secure reference to a particular target experience amongst those candidates. If this picture is correct, then there are systematic connections not only between (1) the degree of purity of a phenomenal concept and (2) the exactness of the knowledge of phenomenal character yielded by that phenomenal concept, but also (3) the degree to which the phenomenal content of that phenomenal concept constrains the candidates for target experiences.²⁷

Conclusion

This paper has argued that knowledge of what it is like to have an experience varies along a spectrum from the more exact to the more approximate. I motivated the degreed picture by appeal to limits in our epistemic abilities. I argued that purity is independent from both determinability and vagueness. I explained how degrees of purity can be understood in terms of the elimination of phenomenal possibilities, where phenomenal concepts that rule out more phenomenal possibilities yield more exact knowledge of what it is like to undergo their target experience. And I discussed the philosophical implications of the degreed picture for questions concerning what it takes to acquire a phenomenal concept, what it takes to master a phenomenal concept, and how phenomenal concepts refer. The result is a richer and sharper picture of what we can know about what it is like. Even if our knowledge of phenomenal character is approximate, our knowledge of phenomenal knowledge may grow increasingly exact.

²⁷ Whether this picture is plausible turns on whether there can be mismatches between what the target experience is like and what the target experience is represented as being like. This depends on some tricky questions about whether alleged cases of mismatch are really cases of reference failure (perhaps akin to concepts with inconsistent contents, such as SQUARE CIRCLE). If it turns out that there can be mismatches of this kind, then strictly speaking not all phenomenal concepts yield *knowledge* of what it is like to undergo their target experiences. But even if that is the case, the degreed picture still captures how one can *think* more or less exactly about phenomenal character.

REFERENCES

- Alter, Torin & Walter, Sven (eds.) (2006). *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*. Oxford University Press. Alter, Torin (2013). Social Externalism and the Knowledge Argument. *Mind* 122 (486).
- Ball, Derek (2009). There are no phenomenal concepts. *Mind* 118 (472):935-962.
- Ball, Derek (2013). Consciousness and Conceptual Mastery. *Mind* 122 (486):fzt075.
- Balog, Katalin 1999: 'Conceivability, Possibility, and the Mind-Body Problem'. *Philosophical Review*, 108, pp. 497-528
- Balog, Kati (2009). Phenomenal Concepts. In Brian McLaughlin, Ansgar Beckermann & Sven Walter (eds.), *The Oxford Handbook of Philosophy of Mind*. Oxford University Press.
- Balog, Katalin, 2012, 'In Defense of the Phenomenal Concept Strategy', *Philosophy and Phenomenological Research*, 84: 1-23.
- Broi, Antonin (2020). Revelation and Phenomenal Relations. *Philosophical Quarterly* 70 (278):22-42.
- Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy*, IV, 73-122.
- Carruthers, P. 2003. Phenomenal concepts and higher-order experiences. *Philosophy and Phenomenological Research*.
- Carter, Sam (2019). Higher order ignorance inside the margins. *Philosophical Studies* 176 (7):1789-1806.
- Chalmers, David J. (2003). The content and epistemology of phenomenal belief. In Quentin Smith & Aleksandar Jokic (eds.), *Consciousness: New Philosophical Perspectives*. Oxford University Press. pp. 220--72.

Chalmers, D., 2007, "Phenomenal Concepts and the Explanatory Gap." in T Alter and S Walter (eds) *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*, Oxford University Press, 167-195.

Chalmers, David (2009). The Two-Dimensional Argument Against Materialism. In Brian P. McLaughlin & Sven Walter (eds.), *Oxford Handbook to the Philosophy of Mind*. Oxford University Press.

Goff, Philip (2011). A posteriori physicalists get our phenomenal concepts wrong. *Australasian Journal of Philosophy* 89 (2):191 - 209.

Goff, Philip (2015). Real acquaintance and physicalism. In Paul Coates & Sam Coleman (eds.), *Phenomenal Qualities: Sense, Perception and Consciousness*. Oxford University Press.

Hill, C.S. 1997. Imaginability, conceivability, possibility, and the mind-body problem. *Philosophical Studies* 87:61-85.

Hill, C.S. & McLaughlin, B.P. 1998. There are fewer things in reality than are dreamt of in Chalmers' philosophy. *Philosophy and Phenomenological Research*.

Horgan, Terry & Kriegel, Uriah (2007). Phenomenal epistemology: What is consciousness that we may know it so well? *Philosophical Issues* 17 (1):123-144.

Jackson, Frank (1982). Epiphenomenal qualia. *Philosophical Quarterly* 32 (April):127-136.

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

Levin, Janet (2006). What is a phenomenal concept? In Torin Alter & Sven Walter (eds.), *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*. Oxford University Press.

Loar, Brian. 1990. Phenomenal states. *Philosophical Perspectives* 4:81-108.

Mahtani, Anna (2008). Williamson on inexact knowledge. *Philosophical Studies* 139 (2):171 - 180.

Margolis, Eric & Laurence, Stephen (2007). The ontology of concepts: Abstract objects or mental representations? *Noûs* 41 (4):561-593.

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

McLaughlin, Brian P. (2012). Phenomenal Concepts and the Defense of Materialism. *Philosophy and Phenomenological Research* 84 (1):206-214.

Michaelson, Eliot and Reimer, Marga, "Reference", *The Stanford Encyclopedia of Philosophy* (Spring 2019 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2019/entries/reference/>>.

Nida-Rümelin, Martine and O Conaill, Donnchadh, "Qualia: The Knowledge Argument", *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2019/entries/qualia-knowledge/>>.

Papineau, D. 2002. *Thinking about Consciousness*. Oxford University Press.

Papineau, David (2006). Phenomenal and perceptual concepts. In Torin Alter & Sven Walter (eds.), *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*. Oxford University Press. pp. 111--144.

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

Pereboom, D., 2011, *Consciousness and the Prospects of Physicalism*, Oxford: Oxford University Press.

Perry, John 2001: *Knowledge, Possibility, and Consciousness*. Cambridge, MA: The MIT Press.

Rabin, Gabriel (2011). Conceptual mastery and the knowledge argument. *Philosophical Studies* 154 (1):125-147.

Rabin, Gabriel Oak (forthcoming). Toward a Theory of Concept Mastery: The Recognition View. *Erkenntnis*:1-22.

Rabin, Gabriel (manuscript). How to Twin-Earth a Phenomenal Concept.

Schroer, Robert (2010). Where's the Beef? Phenomenal Concepts as Both Demonstrative and Substantial. *Australasian Journal of Philosophy* 88 (3):505-522.

Schroer, Robert (2013). Do the Primary and Secondary Intensions of Phenomenal Concepts Coincide in all Worlds? *Dialectica* 67 (4):561-577.

Schroeter, Laura, "Two-Dimensional Semantics", *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2017/entries/two-dimensional-semantics/>>.

Schwitzgebel, Eric, "Introspection", *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), Edward N. Zalta (ed.), forthcoming URL = <<https://plato.stanford.edu/archives/win2019/entries/introspection/>>.

Siewert, Charles, "Consciousness and Intentionality", *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2017/entries/consciousness-intentionality/>>.

Stoljar, Daniel, 2005, 'Physicalism and Phenomenal Concepts', *Mind and Language*, 20: 469–94.

Sturgeon, Scott (1994). The Epistemic View of Subjectivity. *Journal of Philosophy* 91 (5):221-235.

Sundström, Pär (2011). On Imagism about Phenomenal Thought. *Philosophical Review* 120 (1):43-95.

Sundström, Pär (2011). Phenomenal Concepts. *Philosophy Compass* 6 (4):267-281.

Tao, Terence (2011). *An Introduction to Measure Theory Vol. 126*. Amer Mathematical Society.

Tye, Michael (1995). *Ten Problems of Consciousness*. Cambridge, MA, The MIT Press.

Tye, M. 2003. A theory of phenomenal concepts. In (A. O'Hear, eds) *Minds and Persons*. Cambridge University Press.

Weiskopf, Daniel A. (2009). Atomism, pluralism, and conceptual content. *Philosophy and Phenomenological Research* 79 (1):131-163.

Williamson, Timothy (1992). Inexact knowledge. *Mind* 101 (402):217-242.