Knowing What It's Like

Blinded Submission for the *Journal of Philosophy*

WORD COUNT (including footnotes): 9970
KEYWORDS: phenomenal knowledge, phenomenal concepts, concept mastery, inexact knowledge, experience requirement, pure phenomenal concepts

ABSTRACT:

This paper argues that (1) knowledge of what it's like varies along a spectrum from the more exact to the more approximate, and that (2) phenomenal concepts vary along a spectrum in how precisely they characterize what it's like to undergo their target experiences. This degreed picture contrasts with the standard all-or-nothing picture of phenomenal concepts and phenomenal knowledge, where one either possesses a phenomenal concept of an experience or not, and where one either knows what an experience is like or not. I motivate the degreed picture by appeal to limits in epistemic abilities such as recognition, imagination, and inference. Under this view, phenomenal concepts that yield more exact phenomenal knowledge are those that eliminate more possibilities for what it's like to undergo their target experiences. I also argue that approximate knowledge of what it's like cannot be explained merely by appeal to determinable or vague phenomenal concepts, discuss how the degreed picture aligns with the semantics of 'knows what it's like' expressions, and explain how the view challenges some standard assumptions about the acquisition conditions, requirements for mastery, and referential mechanisms of phenomenal concepts.

Introduction

Consider what it's like to feel pain, to see red, or to smell cinnamon. Then consider what it's like to undergo the echolocation experiences of bats, the proprioceptive experiences of octopuses, or the electromagnetic experiences of aliens. There's 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's like to feel pain, see red, and smell cinnamon because you have phenomenal concepts for those experiences. But you don't know what it's 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 that (1) knowledge of what it's like varies along a spectrum from the more approximate to the more exact, and that (2) phenomenal concepts vary along a spectrum in how precisely they characterize what it's 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 standard 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 what it's like cannot be explained merely by appeal to determinable or vague phenomenal concepts; **§4** develops a framework for characterizing the epistemic structure of inexact phenomenal knowledge; **§5** examines the semantics of 'knows what it's like' expressions'; **§6** explains how the degreed picture has substantive consequences for 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's like to have that experience (or, synonymously, the

phenomenal character of that experience). If you think about what it's like to see red, feel pain, or smell cinnamon, then you are deploying phenomenal concepts.[1]

A terminological note: I'll follow the standard convention of using expressions of the form 'what it's like to experience *x*' as reference-fixers (read *de re*) rather than as individuation conditions (read *de dicto*). For example, I'll use the phrase 'what it's like to see red' to denote the kind of experience normal humans in fact have when seeing red objects under normal conditions. On this way of talking, if a subject hallucinates red, then they might still acquire knowledge of what it's like to see red, even if they have never actually veridically perceived red.[2]

Let's say 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 within the extension of A. For example, scarlet is in the extension of RED, but RED isn't a concept of scarlet (it's a concept of red). In general, I'll use the term 'target experience' as neutral between denoting phenomenal properties (such as redness) vs. denoting particular experiences (such as a particular red experience), and I'll take target experiences to be wholly individuated by their phenomenal characters. To simplify the language, I'll usually omit the term 'phenomenal' when talking about concepts, properties, and experiences.

I'll assume throughout the paper that concepts are *mental representations*. More specifically, I'll 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 way of thinking about concepts is standard in cognitive science and is common amongst philosophers of phenomenal concepts.[3]

---

[1] Note that a phenomenal concept must not only refer to an experience *x*, but must also enable one to think about what it's like to undergo *x*. The concept MY FAVORITE MENTAL STATE might happen to refer to an experience, but that doesn't make it a phenomenal concept.

[2] In the framework developed in Mehta [forthcoming]'s analysis of 'what it's like', my concern is with "hard character," or the sense of 'phenomenal character' prominent in discussions of the hard problem of consciousness.

[3] As examples, Chalmers [2003: 4] takes "concepts to be mental entities [that] are constituents of beliefs…analogous to the way in which words are constituents of sentences," Tye [2011: 302] assumes that "concepts are mental representations deployed in thought, belief, and knowledge," and Balog [2012: 9] says that "concepts are mental representations that are…words of Mentalese."

Under 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).[4] 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'll speak in the language of the mental 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.[5]

By possessing a phenomenal concept, one can acquire knowledge of what it's like to undergo the target experience. I favor the view that knowledge of what it's like is constituted by, rather than merely acquirable from, phenomenal concepts.[6] However, the arguments of this paper are compatible (given some terminological substitutions) with views that instead hold that there is merely a causal (rather than constitutive) relationship between phenomenal concepts and knowledge of what it's like. As a neutral locution, I'll talk of phenomenal concepts *yielding* knowledge of what it's like, by which I mean possession of the phenomenal concept puts one in a position to acquire knowledge of what it's like to undergo the target experience.

According to what I'll call 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's like to undergo its

[4] Some readers might think that I'm using 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.

[5] 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.

[6] 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 what it's like even for experiences one has never had (a point I argue for in §3).

target experience.[7] Equivalently, in my view, what it is for a phenomenal concept to have a higher degree of purity is for it to more precisely characterize what it's 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's like to undergo x.

By contrast, according to the standard *all-or-nothing picture*, for any experience x, either you possess a phenomenal concept of x or not: if so, then you know what x is like, and if not, then you don't. Whereas the degreed theorist posits many phenomenal concepts (at varying degrees of purity) for x, the all-or-nothing theorist posits only one phenomenal concept for x (where one possesses that phenomenal concept just in case one can think about what x is like).

Why think that the all-or-nothing picture is standard? 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'll describe. 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.[8]

---

[7] 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 purity are distinct from non-phenomenal concepts of experiences (see §2), and (3) no phenomenal concepts refer directly via their phenomenal character (see §3).

[8] 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 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 Connaill

A second source of evidence is that philosophers sometimes explicitly state that phenomenal concepts yield exact knowledge of what it's like. 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." Notice that in this example, the target experience is a determinable (namely, 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'll 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.[9] It may not yet be obvious why the degreed picture is in tension with these claims; I'll discuss these points in §6.

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 (because phenomenal concepts don't reveal the physical nature of their target experiences).[10] However, the all-or-nothing picture concerns only what phenomenal concepts reveal about *what it's like to undergo* their target experiences. A physicalist may very

---

[2019], and the papers in Alter & Walter [2006]. For discussions of inexact knowledge, see Williamson [1992], Mahtani [2008] and Carter [2019].

[9] Other common commitments that are arguably in tension with the degreed picture, but which I won't discuss in detail, 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]).

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

well hold that phenomenal concepts don't reveal the physical nature of their target experiences yet still assume that phenomenal concepts yield exact knowledge of what it's like. And 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 what it's like, 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 that may be plausible, I think it's quite clear that this hypothesis doesn't 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 hasn't 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's like.

§2 | Motivations

To illustrate the plausibility of the degreed picture, I'll discuss three cases involving subjects whose phenomenal concepts arguably yield only approximate knowledge of what it's like. Each case will concern *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's also stipulate that the property of being a scarlet experience is maximally determinate and that there are no borderline cases of scarlet experience.

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 the color experiences induced by unlabeled color chips. Before day *n*, her ability to recognize scarlet experiences is basically 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 vermillion). 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,[11] 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's natural to think that on day *n* Ms. Scarlet knows only approximately what it's like to see scarlet, while by day *n*+100 Ms. Scarlet knows exactly what it's 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's 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 just by thinking about what it's 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's like to see scarlet (and imagines it as a kind of chromatic experience), and he can still know on that basis that what it's like to see scarlet is more similar to what it's like to see gray than what it's like to hear a trumpet or feel pain. Though both Mr. Rainbow and Mr. Gray possess concepts of scarlet experience, Mr. Rainbow's concept arguably yields more exact knowledge of what it's 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

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

expert on color experiences but who has never had any visual experiences (Mr. Black doesn't 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's like to see scarlet. However, Mr. Gray arguably has a better grasp than Mr. Black of what it's like to see scarlet (even though neither knows what it's like to see scarlet as well as Mr. Rainbow does). If Mr. Gray has some knowledge of what it's like to see scarlet, then Mr. Gray must be able to think about what it's 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's 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's like to see scarlet. You can reliably recognize instances of scarlet experience when it's presented against dissimilar experiences (such as non-red experiences), you can imagine what it's 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 what it's like to undergo those experiences, and you have a better grasp of what it's like to see scarlet than what it's like to echolocate. Therefore, you have at least some knowledge of what it's like to see scarlet. However, I suspect that you cannot reliably recognize scarlet experience when it's presented against extremely similar red experiences, that you cannot imagine exactly what it's like to experience scarlet (as opposed to other nearby red experiences), that you cannot know simply on the basis of your phenomenal concepts that scarlet experience is exactly as similar in hue to crimson experience as it is to amaranth experience, and that you are in a better position to know what it's like to see scarlet if you are actually seeing scarlet than if you are merely thinking about scarlet experience. Therefore, you don't know exactly what it's 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's like. 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 what it's like. While it's 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's 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's easy to see that these arguments generalize to phenomenal concepts for other experiences as well. As other examples, consider the maximally determinate total experience you had upon first waking up this morning, the complex flavor experience you have when tasting an interesting dish, 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's like to have each of those experiences, given your recognitional, imagination, and inferential abilities. But as before, it's also plausible that you don't know exactly what it's 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 may seem in tension with the idea that thinking about an experience involves undergoing the experience that one is thinking about.[12] 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 what it's like?

Now, few philosophers believe that one literally cannot think about what it's like to undergo an experience unless one is actually undergoing that exact

[12] 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."

experience. You can think about what it's 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's like to undergo 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.[13] Those who accept the existence of direct phenomenal concepts might then argue that direct phenomenal concepts yield exact (rather than merely approximate) knowledge of what it's like. However, the degreed picture doesn't 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.

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 don't 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 don't usually 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 don't 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

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

concepts. But even if we accept that there are such asymmetries (an issue I'll discuss in §6), the current objection still leads to counterintuitive consequences. If you don't possess a phenomenal concept of scarlet experience, then (by definition) you cannot think about what it's like to see scarlet. Yet on the face of it, you've been thinking about what it's like to see scarlet the entire time you've been reading this section. What else might you have been thinking about as you considered the cases discussed earlier? Instead of holding that you read about scarlet experience yet were never thinking about scarlet experience, it's more natural to hold that you thought inexactly about what it's 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're thinking about what it's like to see scarlet, you're really thinking about what it's like to see red and drawing an inference from your non-phenomenal knowledge that scarlet experience is a type of red experience. But while that may be one way of acquiring knowledge that scarlet experience is a type of red experience, it's also plausible that you can simply think about what it's 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*[14] just in case it has borderline cases and *sharp* just in case it's not vague. The concept BALD is vague; the concept PHOTON is sharp.[15] Since

[14] I'll 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's somewhat more cumbersome to do so from within an epistemicist framework.

[15] 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

vague concepts are inexact, it may be tempting to think that purity is just a matter of vagueness. And since vagueness is a familiar phenomenon, this may seem to undermine the interest and novelty of the degreed picture.

Although 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 Ms. Scarlet, 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're as competent in thinking about red experience as one could possibly be, and that you know exactly what it's like to see persimmon. Then your concept of red experience is vague since it has borderline cases yet also maximally pure since it yields maximal knowledge of what it's 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 doesn't 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 | The Structure of Purity

From this point forward, I'll take the degreed picture for granted. This section develops a framework that systematizes the relationship between inexact

---

term 'scarlet experience' is more or less pure than the term 'red experience'. This is further evidence that purity is distinct from vagueness.

knowledge of what it's like and degrees of purity of phenomenal concepts. The core idea is that all phenomenal concepts rule out some (and leave open other) "phenomenal possibilities." The more phenomenal possibilities ruled out by a phenomenal concept, the more exact the knowledge of what it's like yielded by that phenomenal concept. A noteworthy benefit is that the framework will enable us to formally disentangle purity from determinability and vagueness.

Let a *phenomenal possibility* be a candidate for what it might be like to undergo a target experience. Under the degreed picture, most phenomenal concepts rule out some (and leave open other) phenomenal possibilities. For example, your phenomenal concept SCARLET rules out the possibility that what it's like to experience scarlet is what it's (in fact) like to experience pain, but (given the arguments from earlier) it doesn't rule out the possibility that what it's like to experience scarlet is what it's (in fact) like to experience crimson.

There's an interesting question of how to best analyze the notion of a phenomenal possibility. Just as we can ask which kind of entity best satisfies the theoretical roles associated with the notion of a possible world, we can also ask which kind of entity best satisfies the theoretical roles associated with the notion of a phenomenal possibility.[16] For our purposes, though, it will suffice to simply think of phenomenal possibilities as phenomenal properties. Here's the reasoning. Phenomenal possibilities are candidates for what it's like to undergo target experiences. Every target experience is either a particular experience or a phenomenal property. In the former case, what the target experience is like will be fully characterized by the maximally determinate phenomenal property characterizing that target experience. In the latter case, the target experience simply is a phenomenal property. Therefore, what the target experience is like can be fully specified by some phenomenal property. Given this, I'll assume henceforth that the set of phenomenal possibilities is identical to the set of all phenomenal properties.

[16] I suspect 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 be modeled by degrees of set membership for possible experiences. However, for limits of space, I won't discuss this approach in detail.

The degree of purity of a phenomenal concept can be specified as the proportion of the set of all phenomenal possibilities that is ruled out by that phenomenal concept.[17] 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.[18] For the rest of the paper, I'll 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 RED$_1$ 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's 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's 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 the target experience is like. A core claim of this paper is that there exist partially pure phenomenal concepts.

A benefit of this framework is that it enables us to 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

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

[18] For notational convenience, assume that we scale purity values logarithmically, so that a purity value of ½ denotes a moderately pure phenomenal concept (rather than an extremely impure phenomenal concept).

represents a phenomenal possibility, where the size of the circle represents degree of determinability and fuzziness represents borderline cases:

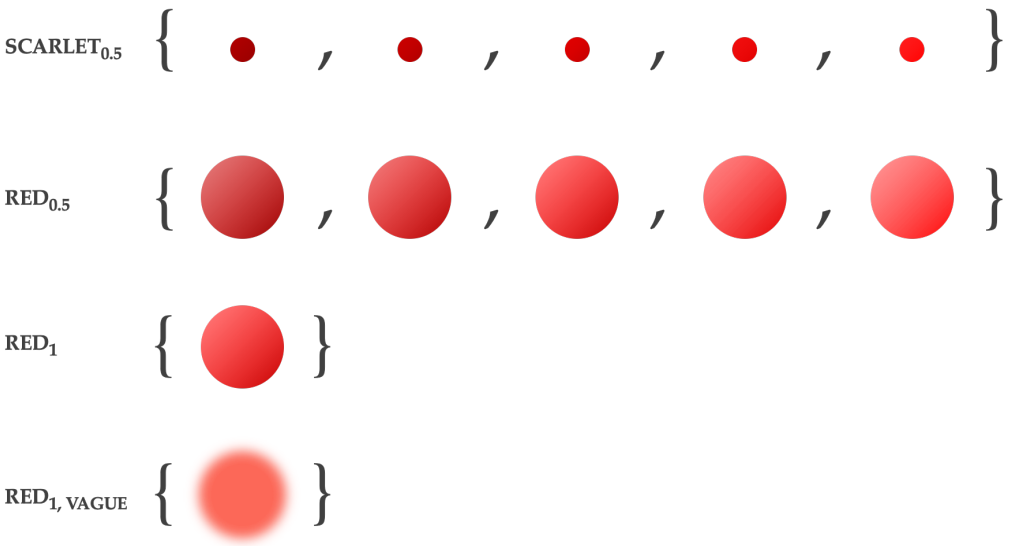

Purity vs. determinability vs. vagueness.

In brief: SCARLET0.5 and RED0.5 differ in determinability (but not purity or vagueness), RED0.5 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 RED0.5 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 RED1 leaves open only one phenomenal possibility while RED0.5 leaves open five, meaning RED1 is purer than RED0.5.

The formal structure of this framework mirrors the epistemic relations between phenomenal concepts. I won't discuss technical details here, but I'll briefly mention a few results. Let A and B be two phenomenal concepts. If A and B leave open exactly the same phenomenal possibilities, then they yield the exact same knowledge of what it's like to undergo their target experiences. If A and B leave open overlapping sets of phenomenal possibilities, then for all one knows (on the basis of A and B), the concepts might pick out phenomenally identical target experiences. If

A and B leave open disjoint sets of phenomenal possibilities, then anyone who possesses both 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 the set of phenomenal possibilities left open by A is a proper subset of the set of phenomenal possibilities left open by B, then A is *strictly purer* than B, meaning that A yields strictly more exact knowledge than B of what the target experience is like.

§5 | The Semantics of 'Knows What It's Like'

I've focused so far on the epistemic structure of phenomenal concepts and phenomenal knowledge. Another resource for investigating knowledge of what it's like is semantic analyses of 'knows what it's like' expressions. How does the degreed picture fare when we look at the semantics?

The standard semantic analyses of 'knows what it's like' expressions take such expressions to be a species of *know-wh* expressions,[19] or expressions where the complement of 'know' is an interrogative clause (headed by 'why', 'when', 'where', 'whether', or 'how') rather than a declarative clause (headed by 'that'). As examples of know-wh expressions, consider 'S knows where the party is' or 'S knows when it's time to quit'. It's clear that the interrogative clauses in these expressions can take on gradable adverbs like 'exactly' or 'approximately': one can know *exactly* where the party is, or *approximately* when it's time to quit. As we have seen over the course of this paper, these adverbial modifiers apply analogously to sentences that attribute knowledge of what it's like:

(1) Ms. Scarlet knows exactly what it's like to see scarlet.
(2) I know approximately what it's like to taste vegemite.

According to standard theories of know-wh expressions, a know-wh sentence is true just in case the subject of the sentence knows an answer to the relevant

[19] See Lycan [1996], Hellie [2004], Tye [2011], Stoljar [2016], and Cath [2019]. The most developed amongst these accounts is Stoljar [2016], who argues that 'knows what it's like' expressions quantify over ways of being affected by events, analogous to how 'knows where' expressions quantify over locations and 'knows when' expressions quantify over times.

embedded wh-question.[20] For example, the sentence 'I know when it's time to quit' is true just in case there is some time (or situation) *t* such that I know that *t* is when it's time to quit. When the sentence contains an adverbial modifier, the truth-conditions are modified accordingly. For example, the sentence 'I know *exactly* when it's time to quit' is true just in case there is some *exact* time *t* such that I know that *t* is when it's time to quit. Following Stoljar [2016], we might then take a sentence like 'Ms. Scarlet knows exactly what it's like to see scarlet' to be true just in case there is some exact way such that Ms. Scarlet knows that it is that way to see scarlet. This indicates that semantic analyses of know-wh sentences (and their modifications by graded adverbs) apply straightforwardly to 'knows what it's like' expressions.

Now, one might object by noting an ambiguity in the phrase 'knows what *x* is like'. The semantic analysis above assumes an *interrogative reading*, where one knows what *x* is like just in case one knows an answer to the question *What is x like?* But there's also a *free-relative* reading, where one knows what *x* is like just in case one knows the experience *x*.[21] The interrogative reading interprets the relevant knowledge as propositional; the free-relative reading interprets the relevant knowledge as objectual. And while gradable adverbs like 'exactly' or 'approximately' sound fine when applied to the interrogative readings of such expressions, they sound off when applied to the free-relative readings of such expressions.

To develop this point, we can paraphrase the relevant 'knows what it's like' sentences in order to force the free-relative reading. One way to do so is by replacing the ambiguous expression 'what *x* is like' with a substitute expression that doesn't contain an interrogative term, such as 'the feeling of *x*'. By doing that, we induce the free-relative reading, under which the sentence is true just in case the subject knows the relevant experience (rather than knows a proposition that answers the

[20] See Stanley & Williamson [2001], Brogaard [2008, 2009], and Pavese [2017] for some discussions. Note that the above authors are primarily concerned with the kind of knowledge attributed by sentences involving infinitival constructions ('I know how to party'), whereas 'knows what it's like' expressions are ostensibly more analogous to sentences containing finite clause constructions ('I know how the party went'). Note also that while Pavese argues that all knowledge is absolute, her analysis is compatible with the claims of this paper (and, more generally, with the existence of inexact knowledge).

[21] See Habgood-Coote [2018] and Stoljar [2018] for discussions of this distinction. See Farkas [2006] and Bengson & Moffett [2011] for defenses of objectual theories of know-wh.

embedded question). When we do this, we find that the aforementioned adverbial modifications no longer work:

? (3) Ms. Scarlet knows exactly the feeling of scarlet.
? (4) I know approximately the feeling of vegemite.

These sentences sound peculiar. But I suspect the peculiarity is merely due to the fact that different kinds of adverbial modifiers apply to interrogative clauses (as in 3) vs. objectual clauses (as in 4). In particular, while some graded adverbs (such as 'exactly' and 'approximately') are inapplicable to objectual clauses, other graded adverbs (such as 'very well' or 'barely') work fine for such expressions. For example, while the sentence 'S knows approximately Paris' sounds odd, the sentence 'S knows Paris very well' sounds fine. This means we can still construct felicitous degreed modifications that apply to the free-relative reading of 'knows what it's like'. As before, let's substitute in 'the feeling of *x*' in order to induce the free-relative reading:

(5) Mr. Scarlet knows the feeling of scarlet very well.
(6) I barely know the feeling of vegemite.

Hence, whether 'knows what it's like' expressions are interpreted in the interrogative or the free-relative sense, the semantics of such expressions aligns with the degreed picture. In fact, this point is strengthened when we consider comparative constructions, such as expressions of the form 'A knows ϕ better than B does' For example, one can say 'A knows better than B where the party is' (interrogative reading) or one can say 'A knows Paris better than B' (free-relative reading). These comparative constructions work just as well for 'knows what it's like' expressions, even when we specifically induce either the interrogative reading (as in 8) or the free-relative reading (as in 9):

(7) Mr. Rainbow knows what it's like to see scarlet better than Mr. Gray does.
(8) Mr. Rainbow knows the answer to the question of what it's like to see scarlet better than Mr. Gray does.
(9) Mr. Rainbow knows the feeling of scarlet better than Mr. Gray does.

These considerations indicate that the degreed picture aligns with standard semantic analyses of 'knows what it's like' expressions. But what is the significance of all this? Well, it's possible that the semantic structure of our 'knows what it's like' expressions mismatches the metaphysical structure of knowledge of what it's like. Therefore, one could acknowledge the observations above yet still reject the degreed picture. But this means that those who reject the degreed picture face an additional explanatory burden. If knowing what it's like is all-or-nothing, then why do our expressions that attribute knowledge of what it's like take on graded adverbs and comparative modifiers? There are ways for the all-or-nothing theorist to resist, but I won't explore the strategies for doing so here.

§6 | Philosophical Implications

The rest of this paper discusses implications of the degreed picture for questions concerning the acquisition conditions, the requirements for mastery, and the referential mechanisms of phenomenal concepts. These discussions will be brief. Their purpose is to exhibit some of the ways in which the degreed picture is philosophically consequential, and to further clarify how the degreed picture works.

Concept Acquisition

This paper began 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. I've argued that this standard explanation oversimplifies.

On the degreed picture, we have relatively exact knowledge of what it's 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's 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 isn't a matter of whether we possess phenomenal concepts for the relevant experiences, but instead a matter of how pure our phenomenal concepts are.

It may strike some as counterintuitive to claim that we possess phenomenal concepts even of bat, octopus, and alien experiences. I suspect that most of that counterintuitive force will be 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's like to have an experience. The least pure phenomenal concepts yield only the knowledge that there's something it's like to have the experience (with no further specificity on what exactly the experience is like). 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's 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 trivializes the acquisition conditions for phenomenal concepts. But notice that there remains a significant difference between concepts that represent experiences *as* experiences (i.e., phenomenal concepts) versus concepts that refer to experiences but don't represent them as experiences (i.e., non-phenomenal concepts that refer to experiences). Though zombies arguably can acquire non-phenomenal concepts that refer to experiences (since it's 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 experiences are like). 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 phenomenal knowledge to maximally

approximate phenomenal knowledge, at which point we cross the threshold to no phenomenal knowledge at all.

A noteworthy consequence concerns the *experience requirement*, or the idea that in order to acquire a phenomenal concept of an experience, one must have undergone that experience (or another experience that is relevantly similar).[22] Against this, I have argued that it's 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 typically 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 doesn't require knowing exactly what it's like to have that experience, the experience requirement is false.

Concept Mastery

I've assumed that concepts are mental representations. By contrast, discussions of concept mastery standardly take concepts to be abstract entities.[23] 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

[22] 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's more theoretically useful and reflective of actual use to define phenomenal concepts as concepts of experiences that enable one to think about what it's like to have those experiences.

[23] See Burge [1979] and Peacocke [1992] for some classic discussions addressing concept mastery. See Rabin [2020] 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.

easy for those of us who favor the mental representation framework to overlook issues about concept mastery.

Under the degreed picture, it's natural to think that concept mastery requires possessing a maximally pure phenomenal concept. If your phenomenal concept of scarlet experience doesn't enable you to know exactly what it's like to see scarlet, then it's plausible that you haven't yet achieved mastery in 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:[24] if one's phenomenal concept doesn't yield exact knowledge of what it's like to have a target experience, then it's plausible that one's phenomenal concept won't permit (or enable) one to endorse (or recognize) the beliefs or inferences diagnostic of mastery. Conversely, it's hard to see where to draw the line if one were to reject maximal purity as a requirement for mastery. It's plausible that at least some knowledge of what it's 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.[25]

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 what it's like at all levels of mastery. In fact, I suspect that part of the reason the degreed picture has been neglected is because philosophers working on phenomenal concepts tend to 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 elide 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.[26]

---

[24] See Rabin [2020].

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

[26] See Nida-Rumelin [2006], 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 such question.

Suppose we grant that the essence of an experience is simply what it's like to undergo that experience (or the set of all phenomenal facts about that experience). Then, if one considers only idealized subjects that have mastery over all their concepts, it is tempting to think that revelation is true. However, once we adopt the degreed picture (and broaden our scope to include subjects who haven't achieved mastery in thinking about experiences), revelation becomes less appealing. Just because one possesses a phenomenal concept of an experience x doesn't mean that one knows exactly what x is like. And if one doesn't know exactly what x is like, then one doesn't know the essence of x. Therefore, while revelation might be true for the special class of maximally pure phenomenal concepts, it's 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's like to undergo its target experience.[27] 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 characters, or as the claim that phenomenal concepts have non-contingent modes of presentation that designate the same referents across intrinsic duplicates.

The degreed picture is in tension with these popular claims about reference. In order for these referential claims to be plausible, it would have to be the case that phenomenal concepts specify exactly what their target experiences are like.[28] In

[27] I'll 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.

[28] 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: 104] says "phenomenal concepts refer to phenomenal properties directly, and not by invoking any further contingent properties of those referents."

other words, this picture of reference works only if all phenomenal concepts are maximally pure. This is because any phenomenal concept that isn't maximally pure leaves open multiple phenomenal possibilities. Since phenomenal possibilities are candidates for what it's like to undergo target experiences, it follows that non-maximally pure phenomenal concepts leave open multiple candidates what it's like to undergo target experiences. By consequence, the phenomenal contents of non-maximally pure phenomenal concepts are insufficient for determining reference to the target experiences of those phenomenal concepts.

Is this plausible? Well, consider again what it's like to see scarlet, and then consider what it's like to see vermillion (which is very similar to scarlet). Speaking for myself, it's not clear there is any difference between how my concept SCARLET EXPERIENCE characterizes what it's like to see scarlet versus how my concept VERMILLION EXPERIENCE characterizes what it's like to see vermillion. If we were to test my recognitional, imaginative, and inferential abilities, it's not clear that such tests would reveal any difference at all in how I think about what it's like to see scarlet versus what it's like to see vermillion.[29] Yet the arguments from §2 indicate that there is good reason to hold that I can think about what it's like to have each experience. So, while it's obvious that SCARLET EXPERIENCE refers to scarlet experience and VERMILLION EXPERIENCE refers to vermillion experience, it's unobvious whether the phenomenal contents of my concepts differ.

In fact, even the phenomenal contents of maximally pure phenomenal concepts arguably don't suffice to 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. Then, on the basis of that concept, one can know exactly what its target experience is like. But there would remain the question of whether the phenomenal concept refers to a particular experience or to a maximally determinate phenomenal property. And even if we were to set aside phenomenal properties and focus only on particular experiences, it's

[29] This doesn't mean the concepts are identical. My claims concern only phenomenal contents, or how the concepts represent *what it's like* to undergo the target experiences. This leaves open the other respects ways in which the concepts may differ: 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.

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.

But if phenomenal contents don't determine reference, 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 natural kind concepts: for example, by definite descriptions, speaker intentions, causal chains, demonstrative applications, deference to experts, rules of use, or some combination of these. There's an interesting question of which of these referential mechanisms are relevant for phenomenal concepts. But since these ideas are familiar, I won't evaluate which of these accounts is most plausible.[30]

Although I've argued that the degreed theorist must abandon the audacious claim that phenomenal contents wholly determine reference to target experiences, the degreed theory remains compatible with the idea that the referential mechanisms of phenomenal concepts differ from the referential mechanisms of other natural kind concepts. A degreed theorist might think that phenomenal contents *constrain* (rather than determine) reference, meaning that the referent of a phenomenal concept must be amongst the phenomenal possibilities the concept leaves open. Suppose, for example, that a phenomenal concept leaves open phenomenal possibilities x, y, and z. Then, according to the present view, even though the phenomenal content of that phenomenal concept doesn't determine which of x, y, or z the concept refers to, it does determine that the referent must be one of x, y, or z. 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.[31]

This means that while the degreed picture undermines the view that phenomenal concepts secure reference merely via the way they represent their target

[30] 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.

[31] 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 (akin to concepts with inconsistent contents, such as SQUARE CIRCLE).

experiences, it remains compatible with a view that is similar in spirit. It's worth noting that if we accept the referential constraint view outlined above, 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 what it's like 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.

Conclusion

This paper has argued that knowledge of what it's 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's like to undergo their target experience. I showed how the degreed picture aligns with standard semantic analyses of 'knows what it's like' expressions. 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's like. Even if our knowledge of phenomenal character remains approximate, our knowledge of phenomenal knowledge may grow increasingly exact.

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.

Brogaard, B. 2008. Knowledge-the and Propositional Attitude Ascriptions, Grazer Philosophische Studien 77/1:147–190.

Brogaard, B. 2009. What Mary Did Phenomenological Research 78/2: 439–467.

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.

Farkas, K. 2016. Know-wh Does not Reduce to Know-That, American Philosophical Quarterly 53/2:109–122.

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.

Habgood-Coote, Joshua (2018). Knowledge-how: Interrogatives and Free Relatives. *Episteme* 15 (2):183-201.

Hellie, B. 2004. Inexpressible Truths and the Allure of the Knowledge Argument, in There's Something About Mary, ed. Y. Nagasawa, P. Ludlow and D. Stoljar, Cambridge, MA: MIT Press: 333–364.

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.

Lycan, W. G. 1996. Consciousness and Experience, Cambridge, MA: MIT Press.

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/>.

Mehta, Neil (forthcoming). The fragmentation of phenomenal character. *Philosophy and Phenomenological Research*.

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, M. (2006). Grasping phenomenal properties. In T. Alter & S. Walter (Eds.), Phenomenal concepts and phenomenal knowledge: New essays on consciousness and physicalism. Oxford: Oxford University Press

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.

Pavese, Carlotta (2017). Know-How and Gradability. *Philosophical Review* 126 (3):345-383.

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 (2020). 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/>.

Stanley, J. and T. Williamson 2001. Knowing How, The Journal of Philosophy 98/8: 411–444.

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

Stoljar, D. 2016. The Semantics of 'What it's like' and the Nature of Consciousness, Mind 125/500: 1161–1198.

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.

Tye, M. 2011. Knowing What It Is Like, in Knowing How: Essays on Knowledge, Mind, and Action, ed. J. Bengson and M. Moffett, New York: Oxford University Press: 300–313.

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.