

# Knowing What It's Like

---

Andrew Y. Lee

Rice University, Dept. of Philosophy

[andrewyuanlee@gmail.com](mailto:andrewyuanlee@gmail.com)

<http://andrewyuanlee.com/>

---

WORD COUNT: 9305 (no footnotes), 11335 (with footnotes)

---

## ABSTRACT:

This paper argues that knowledge of what it is like to have an experience varies along a spectrum from the more exact to the more approximate. I motivate the degreed picture by appeal to limits in epistemic abilities such as recognition, imagination, and inference. I systematize the degreed picture by developing a framework for modeling how phenomenal concepts yield knowledge of phenomenal character. And I explain the implications of the degreed picture for issues about how phenomenal concepts refer, which phenomenal facts are knowable via a phenomenal concept, what it takes to possess a phenomenal concept, and what it takes to achieve mastery in thinking about an experience. A noteworthy consequence is that we can acquire phenomenal concepts even for the exotic experiences of bats, octopuses, and aliens.

---

---

## 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. 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.

The aim of this paper is to argue that the aforementioned asymmetry is a difference in degree, rather than a difference in kind. On my view, knowledge of what it is like to have an experience varies along a spectrum from the more exact to the more approximate. I argue that the degreed structure of phenomenal knowledge is explained by *degrees of purity* of phenomenal concepts: the purer a phenomenal concept, the more exact the knowledge it yields of what it is like to undergo the target experience. I motivate this degreed picture by appeal to limits in epistemic abilities such as recognition, imagination, and inference. I also explain how purity is distinct from both determinability (the relation between determinates and determinables) and vagueness (the possession of borderline cases).

After arguing for the degreed picture, I systematize it by developing a framework for modeling how phenomenal concepts yield knowledge of phenomenal character. The core idea is that all phenomenal concepts rule out *phenomenal possibilities*, or possibilities for what it is like to have the target experience. The more phenomenal possibilities ruled out by a phenomenal concept, the more exact the knowledge of phenomenal character yielded by that phenomenal concept.

From there, I use the framework to explain how the degreed picture challenges standard assumptions about how phenomenal concepts refer, which phenomenal facts one can know on the basis of a phenomenal concept, what it takes to possess a phenomenal concept, and what it takes to achieve mastery in thinking

---

about an experience. A noteworthy consequence of my view is that we can possess phenomenal concepts even for the exotic experiences of bats, octopuses, and aliens.

This paper has three sections: §1 argues for the degreed picture, §2 develops a framework for modeling how phenomenal concepts yield knowledge of phenomenal character, and §3 explains philosophical implications of the degreed picture.

## § 1 | THE DEGREED PICTURE

A *phenomenal concept* is a concept of an experience that enables one to think about what it is like to have that experience. The *target experience* of a phenomenal concept is the experience that the phenomenal concept refers to.<sup>1</sup> By possessing a phenomenal concept, one can acquire *knowledge of phenomenal character*, or knowledge of what it is like to undergo the target experience.

What exactly is a concept? I will take for granted the *mental representation* framework, according to which concepts are mental representations that are individuated by their cognitive roles, that are the constituents of thoughts, and that enable epistemic abilities such as recognition, imagination, and inference. The mental representation framework is dominant in cognitive science and is common amongst philosophers working on phenomenal concepts.<sup>2</sup> The main alternative is the *abstract entity* framework, which takes concepts to be Fregean senses that constitute propositions. Though I prefer to characterize concepts as psychological entities, those who favor the abstract entity framework can straightforwardly translate my talk of concepts as talk of mental representations used to grasp concepts.<sup>3</sup>

On the mental representation framework, for any experience  $x$ , there are many concepts that can be used to think about  $x$ . In general, I will use the locution

---

<sup>1</sup> To simplify the language, I will often omit the term 'phenomenal' when denoting concepts, properties, and experiences—for example, I will use terms such as 'red' or 'red experience' rather than 'phenomenal red' or 'phenomenal red experience'.

<sup>2</sup> For example, Chalmers [2003] says "I take concepts to be mental entities on a par with beliefs: they are constituents of beliefs...in a manner loosely analogous to the way in which words are constituents of sentences."

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

---

'concept of  $x$ ' to mean a concept that enables one to think about  $x$  (i.e.,  $x$  is the referent of the concept). Though every phenomenal concept is a concept of an experience, not every concept of an experience is a phenomenal concept. To be a phenomenal concept, the concept must not only refer to an experience, but must also represent what it is like to have that experience. For instance, the concept BEST POSSIBLE MENTAL STATE might refer to an experience without being a phenomenal concept. Since the concern of this paper is with knowledge of what it is like to have an experience, nearly all of the discussion will focus on phenomenal concepts (rather than non-phenomenal concepts of experiences).

By possessing a phenomenal concept of  $x$ , one is in a position to acquire knowledge of what it is like to have  $x$ . But what kind of knowledge is expressed by the phrase 'knowing what it is like'? My own view is that knowledge of what it is like is a kind of objectual knowledge (distinct from both knowledge-that and knowledge-how) that is grounded in (rather than merely acquirable from) phenomenal concepts.<sup>4</sup> In light of this, I take locutions of the form 'knowing what it is like to have  $x$ ' (for some experience  $x$ ) to denote a different kind of knowledge than locutions of the form 'knowing that  $P$ ' (for some phenomenal fact  $P$ ). Nevertheless, the arguments of this paper are compatible with views that hold otherwise. For those who think that what I call knowledge of phenomenal character just is knowledge of phenomenal facts, or for those who think that phenomenal concepts merely enable (rather than ground) knowledge of phenomenal character, my arguments can be straightforwardly reconstrued within these alternative frameworks. As a neutral locution, I will often simply 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 that knowledge.<sup>5</sup>

---

<sup>4</sup> 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 (because 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 will argue for later). For the purposes of this paper, I will simply take for granted the common assumption that knowledge of phenomenal character requires the possession of phenomenal concepts.

<sup>5</sup> On my view, possession of a phenomenal concept both grounds knowledge of phenomenal character and puts one in a position to acquire knowledge of phenomenal facts.

---

According to the *degreed picture*, phenomenal concepts vary with respect to *degree of purity*, or how precisely a phenomenal concept represents what it is like to have its target experience.<sup>6</sup> Phenomenal concepts with higher degrees of purity yield more exact knowledge of what it is like to have their target experiences, while phenomenal concepts with lower degrees of purity yield more approximate knowledge of what it is like to have their target experiences. If the degreed picture is correct, then 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  $x$ .<sup>7</sup> By contrast, on 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 it is like to undergo  $x$ , and if not, then you do not.<sup>8</sup>

What evidence is there that the all-or-nothing picture is in fact standard? A first source of evidence is that most discussions of phenomenal concepts simply distinguish phenomenal concepts from non-phenomenal concepts of experiences, with the implicit implication that there are no further 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 mentioned above).<sup>9</sup> In fact, some philosophers even explicitly state that phenomenal concepts yield exact knowledge of phenomenal character.<sup>10</sup>

---

<sup>6</sup> I take the term 'purity' from Chalmers [2003], who talks of 'pure phenomenal concepts' (and who endorses the all-or-nothing picture).

<sup>7</sup> As far as I can tell, there is nearly no overlap between the literature on inexact knowledge and the literature on phenomenal concepts. See Williamson [1992] for a classic discussion of inexact knowledge, and see Mahtani [2008] and Carter [2019] for more recent discussions.

<sup>8</sup> A good deal of the phenomenal concepts literature focuses on what makes phenomenal concepts distinctive from other kinds of concepts and how this bears on the mind-body problem. By contrast, this paper is largely neutral on these issues. For a limited sample of papers addressing both 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 a broader overview, see Balog [2009].

<sup>9</sup> A notable exception is Schroer [2013], whose arguments lend some additional support to the view developed in this paper.

<sup>10</sup> As an example, Chalmers [2003] says: "When Mary believes *roses cause [phenomenal red] experiences*, or *I am currently having [a phenomenal red] experience*, she thereby excludes all

---

A second source of evidence is that philosophers working on phenomenal concepts commonly endorse claims that are inconsistent with the degreed picture. These include the claims that phenomenal concepts refer to their target experiences directly via their phenomenal character, that phenomenal concepts put one in a position to know all phenomenal facts about the target experience, and that phenomenal concepts can be possessed only by those who have had the relevant target experience.<sup>11</sup> It may not yet be obvious why the degreed picture is in tension with those theses, but I will address these matters in detail in §3.

Those who already find the degreed picture intuitively plausible may be tempted to think that prior discussions of phenomenal concepts were simply idealizations, rather than endorsements of the all-or-nothing picture. In some cases that may be plausible, though I think the evidence mentioned above indicates that this hypothesis is unlikely to apply across the board. But even if we suppose that the degreed picture is what philosophers 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 previous work to develop a more nuanced picture of knowledge of what it is like.

#### CASES

To illustrate the intuitive appeal of the degreed picture, I will present three cases involving subjects whose phenomenal concepts arguably yield only approximate knowledge of phenomenal character. Each case will concern *scarlet experience*, which we can stipulate to be the kind of color experience normal humans have

---

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.”

<sup>11</sup> Some of the most popular theories of phenomenal concepts suggest these conclusions. These include, for example, *recognitional theories* (which take phenomenal concepts to have modes of presentation that are identical to their referents: see Loar [1990], Carruthers [2003], Tye [2003]) and *constitutive theories* (which take phenomenal concepts to be partially constituted by their target experiences: see Papineau [2002] and Balog [2012]).

---

when looking at scarlet color chips under ideal conditions.<sup>12</sup> Let us also stipulate that there is no more specific way of being a scarlet experience (so scarlet is maximally determinate, at least with respect to color) and that there are no borderline cases of scarlet experience (so SCARLET is not vague). In what follows, I will discuss cases where a subject possesses a phenomenal concept of scarlet experience yet do not know exactly what it is like to see scarlet.

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 the 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 extremely poor. On day  $n$ , 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, Ms. Scarlet acquires a phenomenal concept of scarlet experience the very first time she leaves her room and sees the scarlet color chip.<sup>13</sup> But what explains the changes in her epistemic abilities on the subsequent days? 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

---

<sup>12</sup> The definite description should be interpreted *de re*: what it is for an experience to be a scarlet experience is for it to have a certain phenomenal character.

<sup>13</sup> I say 'conventional wisdom' because of standard views about Mary from Jackson [1982]'s knowledge argument, whose situation parallels the situation of Ms. Scarlet (up until day  $n$ ).

---

merely by whether or not Ms. Scarlet possesses a phenomenal concept of scarlet experience. Instead, it is plausible that from day  $n$  to day  $n+100$ , Ms. Scarlet's phenomenal concept of scarlet experience changes in such a way as to yield more and more exact knowledge of what it is like to see scarlet.

CASE 2: Mr. Rainbow and Mr. Gray are both 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 a rich set of phenomenal facts about scarlet just by thinking about what it is like to see scarlet. Mr. Gray, on the other hand, is monochromatically colorblind (and so has never had a scarlet experience): nevertheless, he still imagines scarlet experience as a kind of color experience, he often discusses the subtle differences between different kinds of color experiences (including scarlet) with Mr. Rainbow, and he can still know 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 just by thinking about scarlet experience. Both Mr. Rainbow and Mr. Gray spend their lives thinking about color experiences, including scarlet experience. But Mr. Rainbow's concept of scarlet experience is arguably purer than Mr. Gray's.

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). Intuitively, Mr. Gray has a better grasp of what it is like to see scarlet than Mr. Black. But if neither 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 knows what it is like to see scarlet better than Mr. Black (even though Mr. Gray also does not know what it is like to see scarlet as well as Mr. Rainbow). 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, since 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, Mr. Gray's phenomenal concept of scarlet experience must be less pure than Mr. Rainbow's.

CASE 3: Consider your own knowledge of what it is like to see scarlet. It is plausible that you have at least some knowledge of what it is like to see scarlet.

---

After all, 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 character of the relevant experiences, and you have a better grasp of what it is like to see scarlet than what it is like to echolocate. But it is also plausible that you do not know exactly what it is like to see scarlet. At least, 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 that scarlet experience is exactly as similar (with respect to hue) to crimson experience as it is to amaranth experience simply by thinking about what it is like to undergo those experiences, and you are arguably in a better position for knowing what it is like to see scarlet if you are actually undergoing a scarlet experience than if you are merely thinking about scarlet experience. Consequently, it is plausible that you possess a phenomenal concept of scarlet experience, but that your phenomenal concept is not maximally pure.

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 (rather than exact) knowledge of what it is like. Though 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. As some other examples, consider the total experience you had upon first waking up this morning, the complex flavor experience you have when drinking a good wine,

---

or the exact kind of visual experience you have when looking at a speckled hen.<sup>14</sup> As before, your phenomenal concepts arguably yield some knowledge of what it is like to have each of those experiences, given the powers of 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.

#### OBJECTIONS

The cases above aim to elicit the intuitive appeal of the degreed picture. But some might remain skeptical. In what follows, I respond to a number of objections to the degreed picture.

A first objection appeals to the idea that thinking about an experience involves undergoing the experience that one is thinking about.<sup>15</sup> 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 it might seem unlikely that thinking about what it is like to have an experience can yield only approximate knowledge of phenomenal character. However, few philosophers literally believe that one cannot think about what it is like to undergo an experience unless one is actually undergoing that exact experience. Consider how 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 severe pain or undergoing severe pain, it is obvious which option is better. The more plausible idea is that thinking about an experience requires instantiating another experience (such as a mental image) that resembles (but is phenomenally distinct from) the target experience. But this is compatible with the degreed picture, since

---

<sup>14</sup> Note that we need not take a stance on whether your visual experience is generic or specific, as is debated in the speckled hen literature. Even if you experience the hen merely as having many speckles (as opposed to a determinate number of speckles), it remains plausible that you have limited epistemic abilities with respect to that generic visual experience. See Poston [2011] and Hasan & Fumerton [2019] for further discussion.

<sup>15</sup> 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.”

---

there is no reason to think that an experience that merely resembles a target experience must yield exact knowledge of what it is like to have that target experience.

In response, some might appeal to *direct phenomenal concepts*, or concepts of occurrent experiences that are partially constituted by those occurrent experiences.<sup>16</sup> Even if most phenomenal concepts are not constituted by their target experiences, some might argue that there is a special class of phenomenal concepts that are constituted in such a way and that these phenomenal concepts yield exact (rather than merely approximate) knowledge of phenomenal character. However, I am making only the modest claim that it is possible to think about what it is like to have an experience even if one is not actually undergoing that experience (rather than that no phenomenal concepts are constituted by their target experiences). In other words, the degreed picture is compatible with the existence of direct phenomenal concepts.

A second objection is that purity is simply a matter of *determinability*, or the relation between determinates and determinables. 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. In the next section, I will explain in systematic terms why purity is distinct from determinability. But for now, let me just make some preliminary remarks to pry the two notions apart.

To begin, note that we do not generally impose such demanding conditions on concept possession. Consider, as examples, 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 a certain way, or that weight is an extrinsic property.<sup>17</sup> In response, the objector might argue that there are asymmetries in concept possession between phenomenal concepts and other kinds of concepts. But even if we accept these asymmetries (an issue that will be discussed in §3), the current objection still leads to

---

<sup>16</sup> See Chalmers [2003] and Horgan & Kriegel [2007] for some arguments in favor of direct phenomenal concepts. See Sundström [2011] for some arguments against.

<sup>17</sup> See Ball [2009] for social externalism arguments concerning concepts of experiences.

---

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. After all, what else might you have been thinking about as you considered the cases discussed earlier? It would be odd to contend that even as you read paragraph after paragraph about scarlet experience, you instead thought only about red experience. The more natural diagnosis is that as you read about scarlet experience, you thought imprecisely about what it is like to see scarlet.

Is there a 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 that represents it only as a red experience. But that is compatible with the degreed picture, for it still entails that you possess a phenomenal concept of scarlet experience since your concept refers to scarlet experience and represents (approximately) the phenomenal character of scarlet. 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 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 radically mistaken about the inferential structure of these conscious mental processes, this response overintellectualizes the psychological story.

A third objection is that degree of purity is simply a matter of vagueness.<sup>18</sup> Let us say that a concept is *vague* just in case it has borderline cases and *sharp* just in case it is not vague: the concept BALD is vague, while the concept PHOTON is sharp.<sup>19</sup>

---

<sup>18</sup> For the purposes of this paper, I will assume that vagueness is a semantic (rather than epistemic) phenomenon. This is principally for ease of explication: the epistemicist about vagueness has analogous reasons for disentangling purity and vagueness, though it is more cumbersome to actually do so within an epistemicist framework.

<sup>19</sup> Vagueness is normally understood to be a property of terms, rather than a property of concepts. By contrast, purity is not a property of terms: for example, there is no sense in

---

Since I characterized purity as a matter of how precisely a phenomenal concept represents its target experience, and since vague representations are imprecise, it may be tempting to think that purity is just a matter of vagueness. But whereas vagueness essentially involves borderline cases, purity is independent of borderline cases. Recall that when we defined the term 'scarlet experience', we stipulated that there are no borderline cases of scarlet experience: any color experience is either determinately a scarlet experience or not.<sup>20</sup> If there are no borderline cases of scarlet experience, then any concept of scarlet experience must be sharp, for 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. In particular, 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 we hold vagueness fixed, it follows that purity is distinct from vagueness.<sup>21</sup>

In the next section, it will become evident exactly how purity, determinability, and vagueness are distinct from each other. But here is a preview: Every phenomenal concept rules out phenomenal possibilities, or possibilities for what it is like to undergo the target experience. Purity is a matter of the proportion of

---

which the term 'scarlet experience' is more or less pure than the term 'red experience'. I take this to be a preliminary reason for thinking that purity is distinct from vagueness. But in order to consider the strongest version of the vagueness objection, I will set this aside and assume that vagueness can also be attributed to concepts.

<sup>20</sup> If you think the natural language term 'scarlet' is vague, then we can always stipulate that 'scarlet' is a technical term that by definition involves no borderline cases.

<sup>21</sup> Is there vagueness without impurity? Suppose that persimmon experience is a borderline case of red experience. Suppose you are as competent in thinking about red experience as one could possibly be and that you know exactly what it is like to see persimmon. Then your concept of red experience is vague since it has borderline cases yet also pure since it yields all possible 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 in CASES) but whose phenomenal concept does not even put them in a position 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.

---

phenomenal possibilities ruled out, determinability is a matter of how specific those phenomenal possibilities are, and vagueness is a matter of the extent to which those phenomenal possibilities have borderline cases.

To summarize, I have argued that knowledge of what it is like to have an experience varies along a spectrum from the more exact to the more approximate. The degreed structure of phenomenal knowledge is due to degrees of purity of phenomenal concepts: the purer a phenomenal concept, the more exact the knowledge of phenomenal character yielded by that concept. For the rest of the paper, I will take for granted that the degreed picture is correct. In §2, I develop the degreed picture in a more technically rigorous way. In §3, I discuss philosophical implications of the degreed picture.

## § 2 | THE STRUCTURE OF PURITY

Though I have argued for the degreed picture, it remains open how to think about purity in more systematic terms. This section develops a framework for modeling how phenomenal concepts yield knowledge of phenomenal character, and in doing so clarifies the structure of purity. As a bonus, the framework also reveals the structural differences between purity, determinability, and vagueness.

### PHENOMENAL POSSIBILITIES

The core idea behind the framework is that every phenomenal concept rules out some (and leaves open other) *phenomenal possibilities*, or possibilities for what it is like to have the target experience. 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 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.

What exactly is a phenomenal possibility? The core theoretical role that phenomenal possibilities play is that they are candidates for what it might be like to undergo a target experience. My aim is to show that every difference in the knowledge of phenomenal character yielded by distinct phenomenal concepts can be captured in terms of which phenomenal possibilities are left open by those phenomenal concepts. In other words, the guiding constraint in an analysis of phenomenal possibilities is that the set of phenomenal possibilities should be exactly rich

---

enough to capture any difference in how phenomenal concepts yield knowledge of phenomenal character.

Every phenomenal concept refers either to a particular experience (e.g., the red experience I am having right now) or to a phenomenal property (e.g., phenomenal red).<sup>22</sup> In light of this, it might be tempting to take the set of phenomenal possibilities to simply be the union of the set of phenomenal properties and the set of particular experiences. Since every target experience is either a phenomenal property or a particular experience, this option is guaranteed to distinguish between all candidates for target experiences. However, this option arguably generates more phenomenal possibilities than necessary, since what it is like to undergo any particular experience is identical to what it is like to undergo the maximally determinate phenomenal property characterizing the total phenomenal character of that particular experience.

What if we instead take the set of phenomenal possibilities to be the set of phenomenal properties? Though this option avoids the redundancies mentioned above, it may also generate more phenomenal possibilities than necessary. Suppose phenomenal hue is instantiated always and only with phenomenal saturation. Then phenomenal hue and phenomenal saturation are distinct phenomenal properties, yet they may not correspond to distinct phenomenal possibilities. At least, one might think that a phenomenal concept that merely represents its target experience as instantiating hue does not differ (with respect to how it represents what it is like to have the target experience) from a phenomenal concept that merely represents its target experience as instantiating saturation. If this is right, then phenomenal properties may be more fine-grained than phenomenal possibilities.

The option I favor is to characterize phenomenal possibilities as sets of possible experiences. A *possible experience* is a maximally specific way that a total experience could be (just as a possible world is a maximally specific way that the world could be). Possible experiences are individuated wholly by phenomenal character, so distinct possible experiences necessarily have distinct phenomenal characters. And possible experiences are maximally complete, in that they specify a total

---

<sup>22</sup> Note that I will always use the term 'target experience' to mean the referent of a phenomenal concept, whether that referent is a phenomenal property or a particular experience.

---

experience (rather than just a part or aspect of an experience). Under this option, the set of phenomenal possibilities is the powerset of the set of possible experiences.

This analysis of phenomenal possibilities arguably captures every difference in how phenomenal concepts yield knowledge of phenomenal character. As noted earlier, phenomenal possibilities are candidates for target experiences, and target experiences are always either particular experiences or phenomenal properties. For any particular experience, the corresponding set of possible experiences is just the singleton set containing the possible experience type-identical to the particular experience. For any phenomenal property, the corresponding set of possible experiences is the set of possible experiences that instantiate that phenomenal property. The more determinable a phenomenal property, the larger its corresponding set of possible experiences; the more vague a phenomenal property, the greater the number of possible experiences that are borderline cases. In light of this, the possible-experiences analysis of phenomenal possibilities seems exactly rich enough to satisfy our initial desideratum of distinguishing between all phenomenal possibilities.

#### THE FRAMEWORK

Now we can develop the framework. The basic idea is that we can capture the knowledge of phenomenal character yielded by any phenomenal concept by specifying the phenomenal possibilities left open by that phenomenal concept.

Let us say that the *subjective content*<sup>23</sup> of a phenomenal concept is the set of phenomenal possibilities left open by that phenomenal concept.<sup>24</sup> For any phenomenal concept, its subjective content can be specified as  $\{x_1, x_2, \dots, x_n\}$ , where each  $x_i$  is a phenomenal possibility left open by that phenomenal concept. Since every phenomenal possibility is itself a set of possible experiences, each phenomenal

---

<sup>23</sup> Note that 'subjective content' does not mean content that supervenes on phenomenal character, as the term is sometimes used in discussions of phenomenal intentionality (see Kriegel [2013] and Bourget & Mendelovici [2016]). I use the term 'subjective content' rather than simply 'content' to leave open the possibility that phenomenal concepts may also have other kinds of contents (e.g., descriptive contents).

<sup>24</sup> I will take for granted that concepts have contents. On this way of understanding content, 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.

possibility  $x_i$  can be specified as  $\{e_1, e_2, \dots, e_n\}$ , where each  $e_i$  is a possible experience. Under this analysis, the degree of purity of a phenomenal concept is a matter of the number of phenomenal possibilities ruled out by that phenomenal concept.

To assign numerical values to degrees of purity, let us say the *purity value* of a phenomenal concept is the proportion of the set of all phenomenal possibilities ruled out by that phenomenal concept.<sup>25</sup> As a result, all phenomenal concepts have a purity value from 0–1 inclusive, where higher values indicate higher degrees of purity. For practical convenience, let us stipulate that we scale purity values logarithmically, so that a purity value of 0.5 denotes a moderately pure phenomenal concept (rather than an extremely impure phenomenal concept).<sup>26</sup> Throughout 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<sub>0.2</sub> denotes a relatively impure phenomenal concept of scarlet while RED<sub>1</sub> denotes a maximally pure phenomenal concept of red.

Now we can identify 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 (though 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.<sup>27</sup> Between those 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. Defining these

---

<sup>25</sup> Since there is an infinite number of phenomenal possibilities, we cannot simply take degree of purity to correspond to the cardinality of the associated sets of phenomenal possibilities. Instead, we need a measure that works for infinitary sets. For limits of space, I will set aside this issue, noting that analogous issues arise also for possible worlds accounts of content.

<sup>26</sup> Without logarithmic scaling, a purity value of 0.5 would correspond to leaving open half of all phenomenal possibilities, which presumably ought to count as extremely impure.

<sup>27</sup> Note that zero purity is distinct from lacking a purity value. A phenomenal concept with zero purity still represents the target experience as being such that there is something it is like to have the experience (whereas a concept of an experience that lacks a purity value does not do so). I discuss this point more in §3.

limit cases will prove useful in §3, when we examine some philosophical implications of the degreed picture.

The framework enables us to dissociate purity, determinability, and vagueness. Whereas purity is a matter of the number of phenomenal possibilities ruled out (i.e., the set size of its subjective content), determinability is matter of how specific those phenomenal possibilities are (i.e., the set sizes of the members of its subjective content), and vagueness is a matter of the extent to which those phenomenal possibilities have borderline cases (where borderline cases might be modeled by degrees of set membership<sup>28</sup>). These formal differences between purity, determinability, and vagueness are illustrated in the following diagram (where each circle represents a phenomenal possibility, where the size of the circle represents the number of possible experiences that comprise that phenomenal possibility, and where fuzziness represents borderline cases):

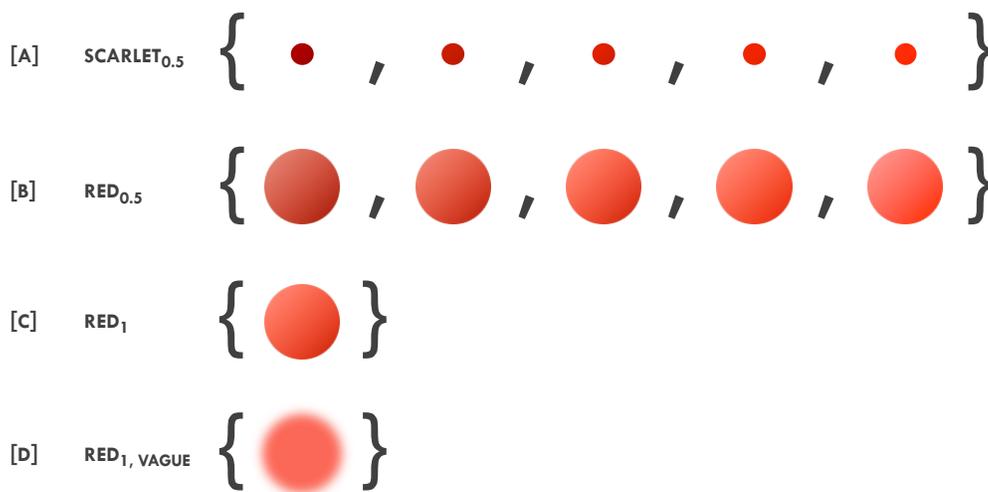


FIGURE 1: Purity vs. determinability vs. vagueness.

In brief: A and B differ in determinability (but not purity or vagueness), B and C differ in purity (but not determinability or vagueness), and C and D differ in

<sup>28</sup> For discussion of fuzzy set theory, see Zadeh [1965].

vagueness (but not purity or determinability). To see how purity dissociates from determinability, consider  $RED_{0.5}$  versus  $RED_1$ . Both phenomenal concepts are equally determinate since they leave open phenomenal possibilities associated with the same number of possible experiences, but  $RED_1$  leaves open only one phenomenal possibility while  $RED_{0.5}$  leaves open five, meaning  $RED_1$  is purer than  $RED_{0.5}$ .<sup>29</sup> To see how purity dissociates from vagueness, consider  $RED_{0.5}$  versus  $RED_1$ . Both concepts are perfectly sharp since neither has phenomenal possibilities with any borderline cases, but  $RED_1$  is purer than  $RED_{0.5}$ .<sup>30</sup>

The set-theoretic structure of the framework is an image of the epistemic structure of phenomenal concepts. To see this, let A and B be two phenomenal concepts.<sup>31</sup> If A and B are associated with exactly the same phenomenal possibilities, then they yield the exact same knowledge of the phenomenal character of their target experiences. If A and B partially overlap, then for all one knows they might pick out target experiences that are phenomenally identical. If A and B are completely disjoint, then one is in a position to know that the target experience of A is phenomenally distinct from the target experience of B. Last, if A is a proper subset of B, then A eliminates all phenomenal possibilities that B eliminates plus more. In such a case, let us say that 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.

---

<sup>29</sup> A noteworthy case is the concept  $EXPERIENCE_1$ , which is maximally determinable (since it applies to every possible experience) yet maximally pure (since the only phenomenal possibility it leaves open is the set of all possible experiences). In my view,  $EXPERIENCE_1$  is a good candidate for being a maximally pure phenomenal concept acquirable by normal humans.

<sup>30</sup> To see how determinability dissociates from purity, consider  $SCARLET_{0.5}$  versus  $RED_{0.5}$ . Both phenomenal concepts are equally pure since they leave open the same number of phenomenal possibilities, but each phenomenal possibility for  $SCARLET_{0.5}$  is more specific than each phenomenal possibility for  $RED_{0.5}$  (meaning that  $SCARLET_{0.5}$  represents its target experience as a more determinate phenomenal property than  $RED_{0.5}$ ). To see how vagueness dissociates from purity, consider the sharp  $RED_1$  concept versus the vague  $RED_1$  concept. Both concepts are equally pure since they leave open only one phenomenal possibility, but the vague  $RED_1$  has borderline cases while the sharp  $RED_1$  does not.

<sup>31</sup> For simplicity, I will ascribe set-theoretic relations to phenomenal concepts, but this should be taken to be elliptical for ascribing those relations to the subjective contents of those phenomenal concepts (since concepts themselves do not stand in set-theoretic relations).

---

### § 3 | PHILOSOPHICAL IMPLICATIONS

In the rest of the paper, I will discuss philosophical implications of the de-greed picture for issues about how phenomenal concepts refer, which phenomenal facts are knowable on the basis of any given phenomenal concept, and what it takes to possess and to master a given phenomenal concept. For limits of space, I cannot discuss any of these implications in detail, but even brief discussions will exhibit some of the wider philosophical significance of the degreed picture.

#### REFERENCE

There is a popular idea that phenomenal concepts refer to their target experiences directly via their phenomenal character. Other ways of expressing this idea are that phenomenal concepts have non-contingent modes of presentation, that the senses of phenomenal concepts suffice to determine their referents, that phenomenal concepts have identical primary and secondary intensions, and that the referents of phenomenal concepts remain constant across internal duplicates.<sup>32</sup> However, the degreed picture is in tension with these claims.

To see why, consider the fact that on the degreed picture, any phenomenal concept that is not maximally pure leaves open multiple phenomenal possibilities. Since phenomenal possibilities are candidates for what it is like to undergo the target experience, that means that these phenomenal concepts leave open multiple candidates what it is like to undergo their target experiences. It follows that the way these phenomenal concepts represent what it is like to undergo their target experiences does not suffice to determine reference to a particular target experience. On the face of it, this seems to contradict the aforementioned claims about direct reference.

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 represents what it is like to see scarlet versus how my concept VERMILLION represents what it is like to see vermillion. If we were to test my recognitional, imaginative, and inferential abilities, it is

---

<sup>32</sup> There are differences between these claims, but they will not matter for the purposes here. See Loar [1990], Chalmers [2003], Tye [2003], Papineau [2007], and Balog [2009] for some well-known expressions of these ideas.

---

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.<sup>33</sup> Yet the arguments from §1 indicate that there is good reason to think that I can still refer both to scarlet experience and to vermillion experience, and I can still think what it is like to have each experience. This means that while it is obvious that SCARLET refers to scarlet experience and VERMILLION refers to vermillion experience, it is unobvious whether the subjective contents of SCARLET and VERMILLION differ at all.

In fact, even the subjective contents of maximally pure phenomenal concepts arguably do not 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: it represents exactly what it is like to undergo its target experience. Nevertheless, there remains the 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 subjective contents of phenomenal concepts invariably underdetermine their target experiences.

Then how do phenomenal concepts manage to refer? In my view, a theory of reference for phenomenal concepts ought to follow a theory of reference for concepts in general. Since there are many theories of reference, there are many options available. As one example, one might appeal to *deference*, where phenomenal concepts refer in virtue of the representational practices of one's community. On such a view, SCARLET refers to scarlet experience in virtue of the fact that the term 'scarlet' is used by color experts to refer to scarlet experience.<sup>34</sup> As another example, one might appeal to *descriptions*, where phenomenal concepts refer in virtue of definite descriptions. For example, when I introduced the term 'scarlet', I stipulated that the

---

<sup>33</sup> This does not mean the concepts have identical cognitive roles. For example, my concept SCARLET (versus VERMILLION) arguably represents scarlet (versus vermillion) experience as the kind of color experience normally caused by scarlet (versus vermillion) color chips. However, such differences are not a matter of the knowledge of phenomenal character yielded by these concepts (and so not differences in subjective contents).

<sup>34</sup> I assume that SCARLET is stipulated to be the concept expressed by the term 'scarlet'.

---

term (and consequently, your concept SCARLET) refers to the maximally specific kind of color experience had by normal humans in idealized conditions. For the purposes of this paper, I will remain neutral on which theory of reference is most plausible, noting only that it is straightforward to appeal to familiar accounts.<sup>35</sup>

My arguments should not be taken to suggest that phenomenal concepts lack any special referential properties. While the degreed picture requires abandoning the idea that subjective contents suffice to determine reference to target experiences, it is compatible with the more modest idea that subjective contents directly constrain the candidates for target experiences. More specifically, one might think that the referent of any phenomenal concept must be amongst the phenomenal possibilities left open by that phenomenal concept. This would allow one to partially retain the claims about reference mentioned at the beginning of the section. Putting it somewhat metaphorically, even though sense and reference do not coincide for phenomenal concepts, they may still overlap. I will remain neutral on whether subjective contents do in fact directly constrain reference,<sup>36</sup> but it is worth seeing that the option is available on the degreed picture.

#### KNOWLEDGE

Towards the beginning of the paper, I mentioned that I take knowledge of phenomenal character (expressed by locutions of the form 'knowing what it is like to have  $x$ ' for some experience  $x$ ) to be distinct from knowledge of phenomenal facts (expressed by locutions of the form 'knowing that  $P$ ' for some phenomenal fact  $P$ ). Though the focus of this paper has been on knowledge of phenomenal character, we can also ask: which phenomenal facts can one know given a phenomenal concept?

On the all-or-nothing picture, the answer is straightforward. If phenomenal concepts always yield exact knowledge of what it is like to undergo their target experiences, then it is natural to think that phenomenal concepts always put one in a position to know all phenomenal facts about their target experiences. After all, there

---

<sup>35</sup> See Michaelson & Reimer [2019] for general discussion of theories of reference.

<sup>36</sup> The issue turns on whether it is possible for there to be a mismatch between what the target experience is like and how the phenomenal character of the target experience is represented. I suspect that at least some potential cases of mismatch may really be cases of reference failure (as in cases of concepts with inconsistent contents, such as SQUARE CIRCLE).

---

would no information about what it is like to have the target experience that is not already contained within the phenomenal concept. But for the discussed in §1, it is implausible that being able to think about what it is like to have that experience suffices for being in a position to know all phenomenal facts about that experience.

Once we recognize the existence of degrees of purity, it is a challenge to explain which phenomenal facts one can know on the basis of any given phenomenal concept. This is because whenever a phenomenal concept yields only approximate knowledge of what it is like to undergo its target experience, that phenomenal concept will yield knowledge of some but not all phenomenal facts about its target experience. For example, your phenomenal concept SCARLET enables you to know that scarlet experience is a type of color experience, but it does not enable you to know that scarlet experience is as similar (with respect to hue) to crimson experience as to amaranth experience. What we need is an account of which phenomenal facts one can know on the basis of any phenomenal concept at any degree of purity.

To begin, observe that for any phenomenal possibility, there is a set of phenomenal properties that applies to all possible experiences that comprise that phenomenal possibility. For example, consider the phenomenal possibility  $x_1 = \{\text{red}_1, \text{red}_2, \dots, \text{red}_{100}\}$ , where each  $\text{red}_i$  is a possible phenomenal red experience. Then the property *red experience* applies to all possible experiences within  $x_1$ , as well as the properties *color experience* and *visual experience*. On the other hand, consider another phenomenal possibility  $x_2 = \{\text{red}_1, \text{red}_2, \dots, \text{red}_{100}, \text{green}_1, \text{green}_2, \dots, \text{green}_{100}\}$ . Then the property *red experience* does not apply to all possible experiences within  $x_2$ , though the properties *color experience* and *visual experience* still do. Since every phenomenal concept is associated with a set of phenomenal possibilities and every phenomenal possibility is associated with a set of possible experiences, we can iterate the procedure above so that we identify the set of phenomenal properties that applies to every possible experience associated with every phenomenal possibility for a phenomenal concept. The result is a set of phenomenal properties that the phenomenal concept's target experience is guaranteed to have.<sup>37</sup>

---

<sup>37</sup> At least, this result is guaranteed if we accept the claim that subjective contents directly constrain reference. However, if subjective contents can misrepresent, then knowledge is conditional on the subjective content being veridical.

This provides the grounds for an account of which phenomenal facts are knowable on the basis of a given phenomenal concept. Suppose a phenomenal concept  $A$  is associated with the set of possible experiences  $\{e_1, e_2, \dots, e_n\}$ . Then for any phenomenal property  $F$  that is instantiated by every possible experience within  $\{e_1, e_2, \dots, e_n\}$ , one can know on the basis of  $A$  that the target experience is  $F$ . As an example, consider a phenomenal concept  $SCARLET_{0.5}$  whose subjective content is  $\{red_1, \dots, red_{100}\}$ . Since all possible experiences associated with  $SCARLET_{0.5}$  are phenomenal red experiences, one can know on the basis of  $SCARLET_{0.5}$  that the target experience of  $SCARLET_{0.5}$  is a type of red experience. However, since the possible experiences associated with  $SCARLET_{0.5}$  include multiple specific kinds of red experiences, one cannot know on the basis of (the subjective content of)  $SCARLET_{0.5}$  which specific kind of red experience  $SCARLET_{0.5}$  refers to.<sup>38</sup>

The account generalizes straightforwardly to phenomenal facts ascribing polyadic properties. Consider the phenomenal fact that red experience is more similar to orange experience than to green experience. This case will involve three phenomenal concepts,  $RED$ ,  $ORANGE$ , and  $GREEN$ , three sets of possible experiences ( $\{r_1, r_2, \dots, r_n\}$ ,  $\{o_1, o_2, \dots, o_n\}$ , and  $\{g_1, g_2, \dots, g_n\}$ ), and the three-place relation  $x$  is more similar to  $y$  than to  $z$ . The account I have sketched predicts that if any  $r_i$  is more similar to any  $o_i$  than to any  $g_i$ , then one can know on the basis of the phenomenal concepts  $RED$ ,  $ORANGE$ , and  $GREEN$  that red experience is more similar to orange experience than to green experience. By contrast, if there is a set of possible experiences  $r_n$ ,  $o_n$ , and  $g_n$  such that it is not the case that  $r_n$  is more similar to  $o_n$  than to  $g_n$ , then one cannot know on the basis of the phenomenal concepts  $RED$ ,  $ORANGE$ , and  $GREEN$  that red experience is more similar to orange experience than to green experience.

To verify that the account is plausible, we can consider its predictions across different kinds of cases. If phenomenal concept  $A$  and phenomenal concept  $B$  are associated with exactly the same phenomenal possibilities, then their subjective contents yield knowledge of exactly the same set of phenomenal facts. If  $A$  and  $B$  overlap,

---

<sup>38</sup> Some might worry about experiences whose instantiation interferes with one's knowledge acquisition processes. Consider, for example, the experience of being bored and distracted. I will assume, however, that we can distinguish how good of an epistemic position one is in from how well one's knowledge acquisition processes function in that position. See Smithies [2019] for more on this distinction.

---

then for all one knows on the basis of A and B, the target experiences of A and B may be the same. If A and B are *disjoint*, then one can rule out the possibility that A and B have the same target experience. And if A is *strictly purer* than B (i.e., A rules out all phenomenal possibilities B rules out and more), then one can know strictly more phenomenal facts on the basis of A than on the basis of B.

We can also consider limit cases. If A is a maximally pure phenomenal concept (i.e., A leaves open a single phenomenal possibility), then one can know all phenomenal facts about A's target experience on the basis of A (since every predicate that applies to the target experience also applies to all phenomenal possibilities for A). On the other hand, if B is a minimally pure phenomenal concept (i.e., B leaves open every phenomenal possibility), then the only phenomenal facts about B's target experience that one can know on the basis of B are those that would hold for any experience whatsoever (since only properties that apply to all experiences apply to all phenomenal possibilities for B).

It may strike some as implausible that knowing a phenomenal fact on the basis of a phenomenal concept requires ruling out all phenomenal possibilities inconsistent with that fact. After all, most philosophers nowadays are fallibilists about knowledge, meaning they think knowledge that P does not require ruling out all possibilities incompatible with P. However, discussions of fallible knowledge tend to focus on empirical knowledge rather than conceptual knowledge, and the arguments for fallible empirical knowledge do not obviously generalize to fallible conceptual knowledge. In fact, canonical examples of conceptual knowledge, such as knowledge of analytic or logical truth, tend to be infallible. Nevertheless, philosophers who favor fallibilism about conceptual knowledge can still take the account above to capture how phenomenal concepts yield knowledge of a core set of phenomenal facts, while also allowing for knowledge of further phenomenal facts beyond that core set.

#### CONCEPT POSSESSION

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 the 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.

This result may lead some to think that the degreed picture is approaching a *reductio ad absurdum*. After all, the claim that we possess phenomenal concepts even of bat, octopus, and alien experiences is counterintuitive. I suspect that much of the counterintuitive force may be residue from implicitly presuming 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 are sufficiently pure, 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 common experiences of normal humans versus our knowledge of the exotic experiences of other kinds of creatures is a matter of degree, and that recognizing this has the diverse implications discussed throughout the paper.

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).<sup>39</sup> These

---

<sup>39</sup> 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

---

minimally pure phenomenal concepts eliminate no phenomenal possibilities (but still represent 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 are led to the result that for any experience, we can acquire at least a minimally pure phenomenal concept of that experience.

This result may raise the worry that the concept of a phenomenal concept has been rendered trivial. But this worry is undercut when we observe that there remains a significant difference between concepts that represent experiences as experiences (i.e., phenomenal concepts) and concepts that refer to experiences but do not represent them as experiences (i.e., non-phenomenal concepts of experiences). Zombies arguably cannot acquire even minimally pure phenomenal concepts, even though they arguably can acquire non-phenomenal concepts of experiences. Though knowing what it is like to have an experience is a matter of degree, there is still a difference in kind between concepts of experiences that represent their referents as experiences and concepts of experiences that do not. 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.

It is often claimed that there is an experience requirement on the acquisition of phenomenal concepts: in order to acquire a phenomenal concept of an experience, one must have had that experience (or another experience that is relevantly similar).<sup>40</sup> But 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 that purity is all-or-nothing, the requirement is less

---

than the difference between B and C. Because of this, I believe my definition of minimally pure phenomenal concepts better carves the conceptual joints.

<sup>40</sup> See, for example, Sturgeon [1994], Tye [1995], and Papineau [2006]. See Ball [2009] for some related arguments against the experience requirement. Since Ball builds the experience requirement into the definition of 'phenomenal concept', he concludes that there are no phenomenal concepts. By contrast, I think it is more apt to define phenomenal concepts as concepts that enable one to think about what it is like to have an experience.

---

compelling once we see that purity comes in degrees. 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 entities.<sup>41</sup> 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. But under the mental representation framework, a change in one's mental representation often means a change in the concept itself. Because of this, issues about concept mastery tend to be elided by those of us who favor the mental representation framework, and it becomes somewhat difficult to even formulate those questions in the first place. In what follows, instead of talk of mastering a concept, I will use the locution 'achieving mastery in thinking about  $x$ '. Then our question is: what does it take to achieve mastery in thinking about an experience?

On the degreed picture, it is natural to think that mastery in thinking about an experience requires possessing a maximally pure phenomenal concept of that experience. If one has a phenomenal concept of scarlet experience that does not enable one to know exactly what it is like to see scarlet, then it is natural to think that one has not 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: 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

---

<sup>41</sup> See Burge [1979] and Peacocke [1992] for some classic discussions addressing issues of concept mastery. See Rabin [forthcoming] for a recent general discussion of concept mastery. For recent discussions on mastery for phenomenal concepts, see Ball [2009, 2013], Rabin [2011], and Alter [2013]. Note that these discussions focus on the ramifications of concept mastery for the knowledge argument, rather than the kinds of issues addressed in this paper.

---

permit (or enable) one to endorse (or recognize) the beliefs or inferences required for mastery.<sup>42</sup> Finally, if one rejects maximal purity as a requirement for mastery, then it is hard to see where to draw the line. 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.<sup>43</sup>

The idea of concept mastery may be 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 neglected is because philosophers working on phenomenal concepts tend to focus on idealized subjects that have already achieved concept mastery. This hypothesis is supported by the fact that the philosophical literature on phenomenal concepts tends to presume that concepts are mental representations (a presumption that tends to obscure issues about concept mastery, as noted above).<sup>44</sup>

A noteworthy upshot is that there is an asymmetry between achieving mastery of phenomenal concepts and achieving mastery of physical concepts. Whereas mastery of phenomenal concepts requires ruling out all phenomenal possibilities but the target experience, mastery of physical concepts arguably does not require ruling out all physical possibilities but the target phenomenon. For example, it is plausible that one can achieve mastery of water concepts even if one does not know that water is H<sub>2</sub>O, whereas I have argued that one cannot achieve mastery of phenomenal concepts unless one knows exactly what the target experience is like.<sup>45</sup>

---

<sup>42</sup> See Rabin [forthcoming] for a recent overview of theories of concept mastery.

<sup>43</sup> Note that I am suggesting possession of a maximally pure phenomenal concept only as a necessary (rather than also sufficient) condition for achieving mastery in thinking about an experience.

<sup>44</sup> Rabin [forthcoming] expresses a related point in more forceful terms: "Theorists of content and of concepts tend to assume, often implicitly, without argument, and without recognizing that any such assumption has been made, that possession of a concept (i.e. the ability to think thoughts containing the concept) entails mastery of that concept. Often they fail to note any distinction at all between possession and mastery."

<sup>45</sup> See Rabin [forthcoming] for further discussion of this asymmetry. The difference may be due to the appearance-reality gap that holds for the physical but not the phenomenal.

---

This asymmetry can explain the appeal and controversy over REVELATION, the thesis that possession of a phenomenal concept puts one in a position to know the essence of its target experience. If REVELATION is true, then anyone who possesses a phenomenal concept of scarlet experience is in a position to know exactly what it is for something to be a scarlet experience. The appeal of REVELATION for phenomenal concepts (but not physical concepts) is explained by the fact that mastering phenomenal concepts requires knowing exactly what it is like to have the target experience<sup>46</sup> (whereas mastering physical concepts arguably does not require knowing the essence of the referent). Because of this, REVELATION rightly identifies an epistemic asymmetry between phenomenal concepts and physical concepts. At the same time, the controversy over REVELATION is explained by the fact that it is false for nearly all phenomenal concepts. Even though maximally pure phenomenal concepts yield exact knowledge of phenomenal character, all other phenomenal concepts yield only approximate knowledge.

#### 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 the degreed structure of phenomenal knowledge is best explained by degrees of purity of phenomenal concepts. I developed a framework for modeling how phenomenal concepts yield knowledge of phenomenal character that systematizes the structure of purity. And I explained the philosophical implications of the degreed picture for questions concerning which phenomenal facts are knowable on the basis of a given phenomenal concept, how phenomenal concepts refer, concept possession, and concept mastery. The result is a richer and sharper picture of the epistemology of experience. In light of this, even if our knowledge of phenomenal character remains forever approximate, our knowledge of the epistemic structure of phenomenal knowledge may grow increasingly exact.

---

<sup>46</sup> I assume the appeal of REVELATION also comes from the idea that knowing exactly what it is like to have an experience enables one to know the essence of that experience.

---

**REFERENCES**

- 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.
- Block, N. & Stalnaker, R. 1999. Conceptual analysis, dualism, and the explanatory gap. *Philosophical Review* 108:1-46.
- Bourget, David & Mendelovici, Angela (2016). Phenomenal Intentionality. *The Stanford Encyclopedia of Philosophy*.
- 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, David J. (2006). Two-dimensional semantics. In E. Lepore & B. Smith (eds.), *The Oxford Handbook to the Philosophy of Language*. Oxford University Press.

---

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.

Gertler, Brie (2015). "Self-Knowledge", *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), Edward N. Zalta (ed.), forthcoming URL = <<https://plato.stanford.edu/archives/win2019/entries/self-knowledge/>>.

Gopnik, A., & Meltzoff, A. (1997). *Words, Thoughts, and Theories*, Cambridge, MA: MIT Press.

Hasan, Ali and Fumerton, Richard, "Knowledge by Acquaintance vs. Description", *The Stanford Encyclopedia of Philosophy* (Summer 2019 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2019/entries/knowledge-acquaindescrip/>>.

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.

Kriegel, Uriah (2013). Two Notions of Mental Representation. In U. Kriegel (ed.), *Current Controversies in Philosophy of Mind*. Routledge. pp. 161-179.

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, William (2019). "Representational Theories of Consciousness", *The Stanford Encyclopedia of Philosophy* (Fall 2019 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2019/entries/consciousness-representational/>>.

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, "Qualia: The Knowledge Argument", *The Stanford Encyclopedia of Philosophy* (Summer 2015 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2015/entries/qualia-knowledge/>>.

Nolan, Daniel (2019). Intensionality and Hyperintensionality. *Routledge Encyclopedia of Philosophy*.

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.

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

---

Poston, Ted, 2007, "Acquaintance and the Problem of the Speckled Hen," *Philosophical Studies*, 132: 331–346.

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.

Rosch, E. (1978). Principles of Categorization, in E. Rosch & B. Lloyd (eds.), *Cognition and Categorization*, Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 27–48.

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/>>.

Smithies, Declan (2019). *The Epistemic Role of Consciousness*. New York, USA: Oxford University Press.

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.

Zadeh, Lotfi A., 1965, "Fuzzy Sets", *Information and Control*, 8(3): 338–353. doi:10.1016/S0019-9958(65)90241-X