

OBJECTIVE PHENOMENOLOGY

ANDREW Y. LEE

NEW YORK UNIVERSITY

ABSTRACT:

This paper develops the foundations for an objective phenomenology—a way of understanding the phenomenal character of experiences that does not require one to have had the kinds of experiences under consideration. My central thesis is that facts that are purely about how experiences are structured are objective. I begin by developing a framework that enables us to assess the idea of objective phenomenology with greater precision. Then I diagnose what makes any given phenomenal fact subjective, argue that there is a class of objective phenomenal facts, and argue for a new kind of explanatory gap.

§ 0 | INTRODUCTION¹

This paper argues for the possibility of an objective phenomenology, or a way of understanding the phenomenal character of experiences that does not require one to have had the kinds of experiences under consideration. My central thesis is that *structural facts* about experience—facts purely about how experiences are structured, as opposed to their particular qualitative characters—are objective.

The idea of an objective phenomenology comes from a cryptic puzzle at the end of the most famous article in contemporary philosophy of mind: “What is it like to be a bat?” by Thomas Nagel. Here is the main passage:

Setting aside temporarily the relation between the mind and the brain, we can pursue a more objective understanding of the mental in its own right ... This should be regarded as a challenge to form new concepts and devise a new method—an objective phenomenology not dependent on [taking up the point of view of the experiential subject] ... Though presumably it would not capture everything, its goal would be to describe, at least in part, the subjective character of experiences in a form comprehensible to beings incapable of having those experiences.

—Thomas Nagel, “What is it like to be a bat?” [1974]

It has been nearly half a century since Nagel’s article was first published, and the passage above is now widely-known. But the idea of an objective phenomenology remains undeveloped. Nagel himself said in the same article that it is “difficult to understand what could be meant by the objective character of an experience.” A common sentiment is that an objective phenomenology, though intriguing, is ultimately impossible.²

¹ Thanks to...

² Since Nagel’s article, there has been little work directly on objective phenomenology. Atkins [2013] discusses the idea of objective phenomenology within a Peircean framework and Mensch [2000] discusses the idea of an objective phenomenology within a Husserlian framework, and Johnston [2007] argues that the contents of minds are objective modes of

In §1, I develop Nagel’s core ideas into a framework that enables us to assess the idea of objective phenomenology with greater precision. In §2, I examine subjective phenomenal facts and diagnose the source of their subjectivity. In §3, I argue that there is a class of objective phenomenal facts: namely, facts purely about how experiences are structured. In §4, I argue for a new kind of explanatory gap, between physical facts and structural facts about experience.

§ 1 | OBJECTIVITY

The basic notion of objective phenomenology is somewhat obscure. To examine the idea rigorously, we need to take Nagel’s core ideas and develop them into a more precise framework. Doing so will enable us to both better understand what an objective phenomenology would be and to better assess its prospects.³

OBJECTIVITY

There are some facts about other creatures that we could understand, such as facts about behavior, function, and physiology. For example, in the case of bats, we could understand facts about flight and feeding behavior, about how their biological systems work, and about the structure of their anatomy. Perhaps more empirical or theoretical investigation is needed to actually acquire knowledge of these facts. But if these facts were presented to us, we would be able to grasp them. These are the kinds of facts that Nagel calls *objective*.

There are also facts about other creatures that it seems we could never understand—specifically, particular facts about what it is like to be those creatures. For example, in the case of bats, these might include facts about what it is like to echolocate.⁴ Our inability to understand such facts is not due to limits in empirical

presentation. However, my aims are quite different from the aims of these other projects.

³ The primary texts from Nagel relevant to these issues are “What is it like to be a bat?” [1974] and *The View From Nowhere* [1986] (especially Chapters 1 and 2).

⁴ Actually, it is likely that we could attain a partial grasp of the qualitative character of echolocation experience, since humans can acquire basic echolocation capacities (even without

or theoretical investigation, but instead due to our inability to occupy the point of view of those creatures. These are the kinds of facts that Nagel calls *subjective*.⁵

The distinction between objective and subjective facts serves as the launching point for Nagel's puzzle. He observes that phenomenal facts seem to necessarily lie on the subjective side. It is hard to see how we could understand facts about what it is like to have bat experiences unless we were to occupy the point of view of a bat, or at least a creature that could have sufficiently similar experiences. The puzzle of objective phenomenology concerns whether there is a way of understanding the phenomenal character of experiences that does not require occupying a particular point of view.

To answer this question, we first need a more precise characterization of objectivity. On the relevant understanding of 'objectivity', a fact is *objective* just in case it is understandable from every point of view. Some paradigm examples of objective facts include mathematical facts, such as $e^{i\pi} + 1 = 0$, and physical facts, such as water is H₂O. Conversely, a fact is *subjective* just in case it is understandable only from particular points of view. Some paradigm examples of subjective facts include facts about what it is like to have particular kinds of experiences, such as phenomenal redness feels like *this* or pain feels like *that*.

Some physicalists might object that phenomenal facts just are physical facts. But for our purposes, 'facts' should be understood to mean true propositions (rather than states of affairs).⁶ We need not assume that distinct facts necessarily correspond to distinct properties or distinct states of affairs. For physicalists who prefer not to frame the discussion in terms of facts, we could instead frame the discussion in terms of truths or concepts. As long as one accepts that there are the aforementioned

training). For an overview of the relevant research, see Kolarik [2014]. Despite this, I will continue using the example of echolocation experience to better align with Nagel's paper.

⁵ Nagel [1974] says that while "there are facts which humans never will possess the requisite concepts to represent or comprehend ... one might also believe that there are facts which could not ever be represented or comprehended by human beings...simply because our structure does not permit us to operate with concepts of the requisite type."

⁶ I choose to use the term 'fact' in part to follow the terminology in Nagel [1974, 1986].

asymmetries between what we could understand about the experiences of other creatures and what we could otherwise understand about the external world, we are in position to examine the puzzle of objective phenomenology.⁷

POINTS OF VIEW

Our initial analysis characterizes a fact as objective just in case it is *understandable from every point of view*. But what does that mean?

On an intuitive level, points of view are meant to capture the different epistemic perspectives associated with different kinds of experiencers. For example, my point of view enables me to understand most phenomenal facts about human experiences but not certain phenomenal facts about bat experiences. On a more precise analysis, *points of view* can be taken to correspond to sets of experiential capacities. Every creature has a set of experiential capacities, which determines which experiences that creature could have. For example, my experiential capacities enable me to have a wide range of human experiences, but they do not enable me to have bat experiences. What individuates my point of view from a bat's point of view is that the experiences that I could have are different from the experiences that a bat could have.

When we apply this to our characterization of objectivity, a fact is objective just in case it is understandable given every set of experiential capacities. More specifically, this concerns every possible point of view, rather than every actual point of view. If we were concerned with every actual point of view, then which creatures are actually instantiated would make a difference to the objectivity of a fact. But intuitively, the objectivity of a fact should not depend upon contingent facts about which creatures exist.

Some might argue that points of view should also be individuated by

⁷ Most physicalists accept that there are such asymmetries. For others, the debate depends on deeper theoretical issues that cannot be addressed here. The issue of objective phenomenology is interesting only if we grant the basic observations expressed by Nagel—I am interested in starting from there in order to progress further.

cognitive capacities, rather than only experiential capacities. But this way of thinking about points of view would fail to capture the concept of objectivity that we started with. For example, a shrimp lacks the cognitive capacities required to understand mathematical facts such as $e^{i\pi} + 1 = 0$ and physical facts such as water is H₂O. But that does not make mathematical and physical facts subjective. If points of view were also individuated by cognitive capacities, then no domain of facts whatsoever would count as objective, since for any fact there would be some points of view that lack the cognitive capacities required to understand that fact. To develop an interesting notion of objectivity in the first place, we should abstract away from cognitive capacities when individuating points of view.⁸

The last key concept is understanding. To *understand* a fact is to grasp its content. So, a fact is *understandable from a point of view* just in case that point of view has the experiential capacities required to grasp the content of that fact. Consequently, a fact might be understandable from a creature's point of view even if that creature lacks the cognitive capacities needed to grasp that fact. For example, suppose my cat lacks the metacognitive capacities needed to think about her own experiences. Nevertheless, facts about her experiences are still understandable from her point of view since she has the experiential capacities required to understand those facts. This analysis also means that a fact might be understandable even if it is unknowable. For example, suppose it is physically impossible to discover that there is an odd number of electrons in the universe. Nevertheless, we would still be able to understand what it is for the universe to contain an odd number of electrons.

To summarize: a fact is objective just in case it is understandable from every point of view, where points of view correspond to possible sets of experiential capacities and where understanding a fact consists in grasping its content. Now we are in position to evaluate our principal question: are there objective phenomenal facts?

⁸ Analogous considerations apply to other kinds of factors that might be used to individuate points of view, such as environmental factors or indexical factors.

§ 2 | QUALITATIVE FACTS

There are indeed objective phenomenal facts—or so I shall argue. But first, I want to diagnose what makes any given phenomenal fact subjective. Identifying the source of subjectivity will set the stage for seeing why some phenomenal facts are objective.

QUALITATIVE FACTS

A *qualitative fact* is a fact that specifies the qualitative character of an experience. Examples of qualitative facts include facts about the feeling of pain, of phenomenal redness, or of echolocation experience. More precisely, a qualitative fact predicates particular qualitative properties to an experience, such as phenomenal redness. Qualitative facts are the kinds of facts we typically have in mind when we think about phenomenal facts. And it is easy to see that canonical examples of qualitative facts are subjective. For example, humans cannot understand certain qualitative facts about the experiences of bats or octopuses or aliens. More empirical investigation or theoretical analysis cannot help in such cases; we are blocked from understanding such facts by the very way that we are built. But is there a deeper explanation of why qualitative facts are subjective?

To understand a phenomenal fact, one must possess the right phenomenal concepts. And to understand qualitative facts about experience, it is plausible that one must possess pure phenomenal concepts. A *pure phenomenal concept* picks out a phenomenal property directly via its phenomenal character.⁹ For example, when I employ the concept PHENOMENAL REDNESS to think about my visual experience of

⁹ For the purposes of this paper, I take for granted that there are pure phenomenal concepts. For an overview and defense of pure phenomenal concepts, see Chalmers [2003 b]. Most philosophers of mind accept that there are such concepts, as evinced by Papineau [2006] when he says that they are “common coin among nearly all contemporary philosophers working on consciousness.” An exception is Ball [2009], who argues that social externalist considerations demonstrate that there are no phenomenal concepts whose acquisition requires one to have had certain experiences, though see Rabin [2011] and Alter [2013] for counterarguments.

a red object, I employ a pure phenomenal concept. In contrast, an *impure phenomenal concept* picks out a phenomenal property via other means, such as definite description or linguistic deference. For example, the concept THE EXPERIENCES BATS USE TO NAVIGATE THEIR ENVIRONMENT is an impure phenomenal concept that picks out echolocation experiences through definite description. While I can acquire impure phenomenal concepts for bat or octopus or alien experiences, I cannot acquire pure phenomenal concepts for those experiences. To acquire a pure phenomenal concept for an experience, it is plausible that one has to have had that experience (or at least a sufficiently similar one, as I discuss below).

In the previous section, I characterized objectivity as a property of facts. But the notion of objectivity can also be naturally extended to concepts. A concept is *objective* just in case it is acquirable from every point of view, and *subjective* just in case it is acquirable only from particular points of view. The observations above mean that pure phenomenal concepts are subjective. And this establishes the first step in our explanation of why some qualitative facts are subjective: a qualitative fact is subjective because pure phenomenal concepts are subjective and because there are limits to which pure phenomenal concepts any given point of view could acquire. This then leads to the next question: what determines which pure phenomenal concepts are acquirable by a given point of view?

CONCEPT ACQUISITION

To answer that question, we must think about the methods that can be used to acquire pure phenomenal concepts. The most obvious method is introspection. For any point of view, there is a core set of phenomenal concepts acquirable through introspection. For example, I might form a pure phenomenal concept of pain by introspecting the pain experiences I have actually had. We need not assume that one can form a pure phenomenal concept of every phenomenal property one's experiences can instantiate. For example, perhaps there are limits in forming pure phenomenal concepts that pick out experiences at the periphery of attention. But it is uncontroversial that one can form pure phenomenal concepts for many of the phenomenal properties one introspects.

Are there also methods that enable one to acquire pure phenomenal concepts for experiences one has never had? I suspect there are two such methods. The first method is *extrapolation*, whereby we form novel concepts characterized by the same dimensions as our prior concepts. For example, even if I have never visually experienced the missing shade of blue, I may be able to extrapolate from my phenomenal concepts of other phenomenal blue experiences to form a phenomenal concept of the missing shade of blue experience. The second method is *recombination*, whereby we recombine concepts for basic experiences we have had into a concept for a more complex experience we have not had. For example, I have never had the experience of eating watermelon while smelling sandalwood. But I have had each experience individually and can form phenomenal concepts for each individual experience. Perhaps by recombining my basic phenomenal concepts, I can acquire a novel phenomenal concept picking out that complex experience.

However, it is plausible there are hard limits in how far introspection, extrapolation, and recombination could take us. Even if a creature had perfect introspective, extrapolatory, and recombinatory capacities, it is unlikely that they would be able to acquire pure phenomenal concepts for experiences radically different from their own. This is evident when we think about the nature of the three methods: introspection is limited to experiences one has had, extrapolation is limited to dimensions of experience whose values one already has pure phenomenal concepts for, and recombination is limited to complex experiences whose constituents one already has pure phenomenal concepts for. But there are plausibly many experiences we have never had that instantiate fundamentally different phenomenal qualities, rather than just different values along the same dimensions or different combinations of the same constituents. Perhaps some of the most exotic experiences of bats or octopuses fall within this category.¹⁰ Consequently, it is plausible that there are principled limits on which pure phenomenal concepts could be acquired by any individual point of view.

¹⁰ This echoes Nagel [1974] when he says that “if extrapolation from our own case is involved in the idea of what it is like to be a bat, the extrapolation must be incomplete.”

This gives us a deeper explanation for the subjectivity of qualitative facts. Understanding qualitative facts requires employing pure phenomenal concepts. But for any point of view, only a limited range of pure phenomenal concepts are acquirable since there are limits to introspection, extrapolation, and recombination. So, any qualitative fact whose understanding requires pure phenomenal concepts acquirable by only certain points of view must be subjective. It is worth noting that this diagnosis still leaves open the possibility that some qualitative facts are objective. Even if each point of view has access to only a limited range of qualitative facts, there might still be some qualitative facts that are accessible from all points of view. In fact, I will argue in §3 that there is at least one kind of qualitative fact that is objective. Nevertheless, we have a general explanation of what makes any given qualitative fact subjective.

DEGREES OF OBJECTIVITY

Before moving on to the objective phenomenal facts, it is worth taking a brief detour to examine degrees of objectivity. Our focus so far has been on *perfect objectivity*, or on whether a fact is understandable from every point of view. But objectivity can also be understood as coming in degrees: a fact is *more objective* when it is understandable from a greater range of points of view.¹¹ For example, qualitative facts about phenomenality itself are intuitively more objective than qualitative facts about red₃₄ experience. This raises a new question: what makes a qualitative fact more or less objective?

Here is a conjecture that seems plausible (but which I will argue is false): degree of objectivity corresponds to degree of *generality*. A phenomenal fact is more *general* when it predicates properties instantiated by a wider range of possible

¹¹ Nagel [1986] also makes the distinction between the binary and graded notions of objectivity when he says, “Though I shall for convenience often speak of two standpoints, the subjective and the objective...the distinction between subjective and objective views is really a matter of degree...” Note that whenever I talk about objectivity without qualification, I have in mind perfect objectivity.

experiences, and more *specific* when it predicates properties instantiated by a narrower range of possible experiences.¹² For example, a maximally general qualitative fact predicates only the phenomenal property of experience itself. In contrast, a maximally specific qualitative fact might predicate the maximally determinate phenomenal property that characterizes the exact total experience I am having right now.

It is natural to think that generality correlates with objectivity. Since general facts predicate qualitative properties instantiable by a wide range of points of view, there are many points of view that have the experiential capacities required for understanding those facts. And since specific facts predicate qualitative properties instantiable by a narrow range of points of view, there are few points of view that have the experiential capacities required for understanding those facts. Hence the conjecture: the more general a fact (the greater the range of experiences that instantiate the properties predicated by the fact), the more objective (the greater the range of points of view from which that fact is understandable).

Yet surprisingly, there are counterexamples to this conjecture. In the next section, I will argue that generality and objectivity come apart when we consider structural facts about experience. Nevertheless, although the conjecture is false for phenomenal facts in general, it is true for qualitative facts. More precisely, there is a link between the specificity of the qualitative properties predicated by a phenomenal fact and the degree of objectivity of that fact. Our conjecture holds true under that restriction, even though it does not hold true simpliciter.¹³

¹² Note that a fact could be general even if it is about specific individuals. For example, the fact that Tom is in pain is highly general because the phenomenal property picked out (namely, painfulness) is instantiated by a wide range of possible experiences.

¹³ This leads to the further question of which kinds of qualitative properties are instantiated by a wider range of possible experiences. This is a hard question that I will set aside here.

§ 3 | STRUCTURAL FACTS

We are now in position to move to the objective phenomenal facts—namely, structural facts. Structural facts are less familiar than qualitative facts within analytic philosophy of mind.¹⁴ But they are interesting and important in their own right—not least because they are objective.¹⁵

STRUCTURAL FACTS

A *structural fact* is a fact purely about the structure of experience. More precisely, structural facts predicate structural properties to experiences without predicating any particular qualitative properties. For example, a structural fact might state that there are three experiences, each of which instantiates certain values along two different dimensions and which stand in certain parthood relations.

This immediately raises the question: what is structure? I am inclined to think that STRUCTURE is a primitive concept—it is hard to think of any concept more general or fundamental.¹⁶ But here are some positive remarks that gesture at the concept: Structure is that which is captured through formal representations, such as mathematical models. Structure is purely about how things relate to each other, rather than what those particular things are. Structure is form, rather than substance. And structure abstracts from the qualitative.

We can also point to examples. Canonical examples of structural concepts include quantitative concepts (e.g., NUMBER), mereological concepts (e.g., PART),

¹⁴ Beyond analytic philosophy, there are historical traditions that have addressed philosophical issues concerning the structure of experience that relate to the issues examined here. These traditions include early psychophysics (e.g., Fechner [1860], phenomenology (e.g., Husserl [1913]), and logical positivism (e.g., Carnap [1928]). For the purposes of space, I will set aside discussing the particular ways in which the ideas discussed here relate to those of authors across these traditions.

¹⁵ Nagel [1974] hints at this idea when he says that “structural features of perception might be...accessible to objective description, even though something would be left out.”

¹⁶ See Sider [2011] for a defense of the claim that STRUCTURE is a primitive concept as well as a more extensive discussion of structure.

logical concepts (e.g., SOME), and nomic concepts (e.g., CAUSE). Other candidates include modal concepts (e.g., NECESSITY), informational concepts (e.g., ENTROPY), and natural concepts (e.g., SIMILARITY). There are also concepts for second-order phenomenal properties such as INTENSITY, CENTRALITY, SALIENCE, and RESOLUTION, though it is less obvious whether these are purely structural.¹⁷ But even within just the clear cases, we have a rich array of structural concepts.

Structural facts predicate both structural properties and the qualitative property of being an experience. This means that to understand a structural fact, one must possess two kinds of concepts. The first consists of the structural concepts mentioned above, such as PART, MAGNITUDE, and SIMILARITY. The second consists of a single pure phenomenal concept: namely, EXPERIENCE. This means that establishing that structural facts are objective requires two arguments: first, that EXPERIENCE is objective, and second, that structural concepts are objective.

OBJECTIVE CONCEPTS

Why does understanding a structural fact require employing the concept EXPERIENCE? Some might contend that if we are interested in structural facts, then we should instead be concerned with *pure structural facts*, or facts that predicate only structural properties. But arguing that pure structural facts are objective would establish nothing about the prospects for objective phenomenology. Our concern is with phenomenal facts about experiences, rather than facts about pure structure. Because of this, the facts under consideration must specify that the objects they refer to are experiences. Without this requirement, it would be unspecified whether the facts we are talking about characterize the structures of experiences, of physical objects, or even of abstract objects.

Is the concept EXPERIENCE objective? Some might think that my previous

¹⁷ It is less clear whether these second-order concepts are purely structural—perhaps they involve a minimal qualitative component. But even if that is so, they would likely play the same kinds of roles as purely structural concepts in actual inquiry into experience. For discussion of centrality structure and salience structure, see Watzl [2017]. For discussion of resolution structure, see Block [2015].

arguments concerning the subjectivity of pure phenomenal concepts are in tension with the objectivity of EXPERIENCE. But the previous section provided a diagnosis of why any given pure phenomenal concept is subjective, rather than an argument that all pure phenomenal concepts must be subjective. The diagnosis was that acquiring pure phenomenal concepts requires introspection, extrapolation, or recombination, and the argument was that one's experiential capacities constrain which pure phenomenal concepts one could acquire. This is consistent with thinking that some pure phenomenal concepts are nevertheless acquirable from all points of view.

In fact, EXPERIENCE is a special case, since it is the maximally general pure phenomenal concept. Since every point of view must have some experiential capacities, there is no point of view that lacks the experiential capacities required to acquire EXPERIENCE. This does not mean that every creature actually has the concept, but it does mean that every creature has the requisite experiential capacities needed to acquire the concept. So even though EXPERIENCE is a pure phenomenal concept, it is nevertheless objective.

What about structural concepts? It is plausible that structural concepts are objective in the same way that mathematical and physical concepts are. Most creatures do not actually have the capacities required to acquire structural concepts, but these are plausibly due to the cognitive limits that prevent them from acquiring certain mathematical and physical concepts, rather than the experiential limits that prevent them from acquiring certain pure phenomenal concepts. So long as we take mathematical and physical concepts to be objective, structural concepts plausibly have the same status.

Some might think that acquiring a structural concept would require one to have had an experience instantiating the corresponding structural property. For example, perhaps acquiring the concept PART requires one to have had an experience that instantiates parthood structure. But the general principle that this inference appeals to is dubious. I can acquire the concepts UNCOUNTABLE INFINITY and IMAGINARY NUMBER, even though it is unlikely that my experiences instantiate the structural properties picked out by those concepts.

What about structural concepts that are qualitative in nature? For example, even though the concept MORE PAINFUL THAN predicates structural properties, it is intuitively subjective. However, when I talk about structural concepts, I mean only structural concepts that have no qualitative components. The concept MORE PAINFUL THAN is not a structural concept in this sense, for it has both a structural component (i.e., GREATER IN MAGNITUDE THAN) and a qualitative component (i.e., PAIN). In contrast, the concept GREATER IN MAGNITUDE THAN, a structural concept, involves no qualitative component. And whereas the former class of concepts are plausibly subjective, the latter class of concepts are plausibly objective.

The upshot is that there is good reason for thinking that both EXPERIENCE and structural concepts are objective. Furthermore, this conclusion is supported by intuitions about cases. While humans cannot understand qualitative facts about bat, octopus, or alien experiences, we can understand structural facts about such experiences. If we were told that the echolocation experiences of bats stand in certain parthood relations to each other, or that the sensory experiences of octopuses have a certain number of dimensions of variation, or that similarity relations between alien experiences are representable by particular quality space structures, then we would be in a position to understand those facts.

Here is a recap of the overall argument of this section: Understanding structural facts requires only structural concepts and the concept EXPERIENCE. Structural concepts are acquirable even by creatures whose experiences do not instantiate the properties picked out by those concepts and are plausibly as objective as mathematical and physical concepts. And EXPERIENCE is the maximally general pure phenomenal concept, which is acquirable from every point of view. Thus, structural facts are objective.¹⁸

¹⁸ Nagel [1974] makes a suggestion in the same spirit when he says that “concepts alternative to those we learn in the first person may enable us to arrive at a kind of understanding even of our own experience which is denied us by the very ease of description and lack of distance that subjective concepts afford.”

SUBSTANTIVITY

Some might agree that structural facts are objective but question their substantivity. Perhaps after we extricate all qualitative content from the facts under consideration, what we are left with is too impoverished to be interesting.¹⁹ Or perhaps what we care about in investigating experience is only knowledge of qualitative facts, rather than knowledge of structural facts.

To see why structural facts are substantive, consider first a specific example: color experiences. We know that human color experiences can be represented using a three-dimensional model, where each point in the space corresponds to a specific color experience, distance in the space corresponds to degree of similarity, and the range of instantiable color experiences forms an asymmetrical shape. By developing such a model, we acquire not only piecemeal knowledge of specific color qualities, but also systematic knowledge of how those color qualities relate to each other. And if we learned that other creatures (such as butterflies or mantis shrimp) have color experiences with different structural properties, then we would enrich our knowledge of those experiences, even if we could not grasp the corresponding qualitative facts.

Second, consider specificity. As a reminder from our earlier discussion, a phenomenal fact is more *general* when it predicates properties instantiated by a wider range of possible experiences, and more *specific* when it predicates properties instantiated by a narrower range of possible experiences. We are now in a position to see why our previous conjecture linking objectivity to generality is false when applied to all phenomenal facts, even though it holds true when restricted to qualitative facts. Some structural facts are highly specific, in that they predicate structural properties that characterize only a narrow range of experiences. But since

¹⁹ This is related to Newman's problem for structural realist theories in philosophy of science: if we understand structure in set-theoretic terms, where a structure consists of a domain of objects and a set of relations on those objects, then all we can glean from structural facts are facts about the cardinality of the objects of a domain. However, the notion of structure discussed here is arguably richer than the notion that leads to Newman's problem. For an overview of Newman's problem, see Ainsworth [2009].

structural facts in general are objective, even the most specific structural facts would still be perfectly objective. Consequently, specific structural facts provide a counterexample to our previous conjecture linking objectivity with generality. This observation also demonstrates in another way why structural facts are substantive: every structural fact with even a minimal degree of specificity provides information about the experiences they refer to, in the sense of eliminating possibilities about what the experience is like. So, structural facts are substantive because they are informationally contentful.

Third, consider a comparison between structural facts about experience and structural facts about the physical world. Most theorists think that much of our knowledge of the physical world consists in knowledge of its structure, with some even arguing that structural facts comprise all of our knowledge of the physical world. Yet almost everyone agrees that our knowledge of the physical world is substantive, rather than impoverished. By the same lights, structural facts about experience are substantive, rather than impoverished.

A final point worth mentioning: for the purposes of carving epistemic joints, I have focused on perfectly objective phenomenal facts. But for the purposes of actual inquiry into experience, we are likely to mostly be concerned with facts that have both structural and qualitative components. Only creatures that are inconceivably exotic (relative to our own point of view) would have experiences for which we would be limited to understanding only structural facts.

PHENOMENALITY

At this point, some might wonder: are structural facts about experience genuinely phenomenal facts? Perhaps you agree with me that structural facts are objective and substantive but think that only qualitative facts truly deserve the label 'phenomenal'. But if structural facts are not phenomenal facts, then I have not made a case for objective phenomenology after all. Of course, this is largely a verbal issue about how we use the term 'phenomenal fact'. Nevertheless, there is good reason to think of structural facts as a species of phenomenal facts.

To start, it is worth distinguishing between structural facts about experiences

from pure structural facts, which predicate only structural properties (and not the qualitative property of being an experience). Pure structural facts are not phenomenal facts, since they could just as well characterize physical or mathematical objects instead of experiences.²⁰ In contrast, structural facts, in the sense I am talking about, predicate not only structural properties, but also the qualitative property of being an experience. On top of that, the structural properties predicated by structural facts describe the phenomenal character of experiences.

Consider again a comparison with structural facts about the physical world. Imagine someone who claims that we hardly learn any physical facts through scientific inquiry because we learn only about the structure of the physical world. That claim seems absurd: it is natural to think that structural facts about the physical world just are a kind of physical fact. By analogy, it is natural to think that structural facts about experience just are a kind of phenomenal fact.

Structural facts also exhibit many canonical properties of phenomenal facts. First, phenomenal facts are the kinds of facts we learn through introspection—and we learn structural facts through introspection. Second, phenomenal facts are the kind of facts that characterize what it is like to have an experience—and structural facts characterize what it is like to have an experience. Of course, structural facts alone cannot completely characterize what it is like to have an experience, but neither can purely qualitative facts. And third—as I argue in the next section—there is an explanatory gap between physical facts and structural facts, just as there is an explanatory gap between physical facts and qualitative facts. A diagnosis is that there is an explanatory gap between physical facts and all phenomenal facts, and that structural facts are a species of phenomenal fact.

The upshot is that structural facts about experience are objective, substantive, and phenomenal. This vindicates Nagel's initial speculation: there is indeed a class of objective phenomenal facts.

²⁰ Analogously, structural concepts are not phenomenal concepts (since they do not attribute any qualitative property) and structural properties are not phenomenal properties (since they do not have any qualitative aspect).

§ 4 | THE STRUCTURAL EXPLANATORY GAP

In this final section, I will argue for a new kind of explanatory gap between physical facts and structural facts about experience. Then I examine, in light of this gap, the implications of objective phenomenology for the investigation of experience.

TWO EXPLANATORY GAPS

When we investigate experience, we use both third-person and first-person methods. *Third-person methods* are perceptual in nature and produce observations from an external point of view. For example, when we examine a subject's experience using perceptual discrimination tasks or brain-imaging technology, we employ third-person methods. In contrast, *first-person methods* are introspective in nature and produce observations from an internal point of view. For example, when a subject attends to their experience in order to identify and analyze its character, they use first-person methods. There is a temptation to think that third-person methods enable the discovery of objective facts and that first-person methods enable the discovery of subjective facts. But structural facts about experience—though objective—cannot be discovered through third-person methods alone.

Consider first the more familiar qualitative explanatory gap. Even if we knew all the relevant third-person facts about a creature, such as facts about brain states and behavior, we would not thereby be able to deduce the qualitative facts about that creature's experience. For example, if we were to build a robot with a sophisticated cognitive architecture, there would be open questions about whether the robot would be conscious at all and what the qualitative character of its experience would be like if so. Since this idea has already been discussed extensively in the philosophical literature on consciousness, I will not discuss it in more detail.²¹

²¹ Most contemporary philosophers of mind accept that there is an explanatory gap (see Levine [1983] and Chalmers [2003 a] for discussion), and for the purposes of this paper I will take the existence of the explanatory gap for granted. Some philosophers think there is an explanatory gap but prefer to characterize the gap in terms besides deducibility (such as in

What is less obvious is that there is also a *structural explanatory gap*.²² Even if we knew all the third-person facts about a creature, we would not thereby be able to deduce structural facts about that creature's experience. This is most obvious when we observe that there is an explanatory gap between third-person facts and facts about which creatures are conscious. Structural facts predicate the property of being an experience. If the third-person facts underdetermine the facts about which creatures are conscious at all, then it trivially follows that the third-person facts underdetermine the structural facts.

However, the structural explanatory gap goes deeper. Some might think that we could deduce structural facts about experiences from the third-person facts if only we knew which creatures are conscious, or if only we knew which particular qualitative properties their experiences instantiate. But even if we had such knowledge, the structural explanatory gap would remain. Suppose that we know all of the third-person facts about the robot mentioned previously, and that we also know all the purely qualitative facts about the robot's experience (including the fact that the robot is conscious in the first place). Suppose also that according to one hypothesis, the robot's experience has a single dimension of variation and no mereological structure, while according to another hypothesis, the robot's experience has multiple dimensions of variation and multiple parts bearing certain similarity relations to each other. Neither of these hypotheses says anything about the qualitative character of the robot's experience; the hypotheses differ only with respect to which structural properties they ascribe. But the two hypotheses, though mutually exclusive, are both consistent with our prior knowledge. This means that knowing which qualitative properties an experience instantiates does not put in a position to deduce how those qualities are structured.

Here are the upshots: Since structural facts are not deducible from third-person facts, there is a structural explanatory gap. Since structural facts are not

terms of entailment, inference, scrutability, or knowledge). Which of these characterizations one favors should not make a difference to my arguments for the structural explanatory gap.

²² The structural explanatory gap is related to the grain problem, which features in Lockwood [1993], and the structural mismatch problem, which features in Chalmers [2017].

deducible from purely qualitative facts, the structural explanatory gap is distinct from the qualitative explanatory gap. On the picture I favor, there is a general phenomenal explanatory gap with two sides: qualitative and structural.

STRUCTURAL BRIDGING PRINCIPLES

What does the structural explanatory gap mean for the investigation of experience? The immediate consequence is that third-person methods alone are inadequate for discovering structural facts about experiences. This raises the question of whether first-person methods would fare better. However, it is easy to see that first-person methods are also limited, since each subject has first-person access only to their own experiences. Even if first-person methods are useful for investigating the structures of our own experiences, they are inadequate for discovering structural facts about the experiences of other creatures.

In my view, the best approach to overcome these limitations is to develop *structural bridging principles*, or principles by which we make inferences about structural facts from physical facts. We might think of such principles as modeled by functions where the input is a set of physical facts (such as facts about a creature's neural state) and the output is a set of structural facts about experience (such as facts about mereological and quantitative structure).

Developing such principles requires combining first-person and third-person methods. Here is the general methodology: By comparing first-person data with third-person data, we establish correlations between physical facts and phenomenal facts. By systematizing and generalizing those correlations, we develop principles about how phenomenal structure maps to physical structure. And by applying those principles to third-person data, we discover structural facts about the experiences of other creatures. More specifically, we can use third-person methods to discover physical facts about a creature, and then apply structural bridging principles to infer phenomenal facts about their experiences.²³ This methodology enables us to discover

²³ See Chalmers [2004] for a more comprehensive discussion of bridging principles and their role in the science of consciousness.

structural facts about the experiences of other creatures, as well as our own.

When we think about the experiences of other creatures, we often tacitly employ structural bridging principles. For example, we might make inferences about the structure of bat experiences from knowledge of bat behavior, function, and physiology. Even the assumption that the structure of an experience must be systematically related to the structure of its physical correlates is a substantive (though extremely plausible) structural bridging principle. But while we implicitly invoke such principles, it is not obvious which particular principles are justified. If we could explicitly characterize the structural bridging principles that we tacitly employ and evaluate their plausibility, then that would be a way of making progress.

What about *qualitative bridging principles*, or principles by which we make inferences about qualitative facts from physical facts? It is obvious that qualitative bridging principles are also vital for the investigation of experience. But they have limitations that do not apply to structural bridging principles. Since qualitative facts are typically subjective, there are limited applications for qualitative bridging principles when we investigate experiences beyond our own point of view. After all, a qualitative bridging principle is useless if we are unable to understand the qualitative facts it outputs. In contrast, since structural facts are objective, no analogous problem arises for structural bridging principles. No matter how exotic the experiences we investigate might be, we would be able to understand structural facts about those experiences. Because of this, structural bridging principles are especially promising for advancing our knowledge of consciousness.

The ultimate goal would be to develop structural bridging principles that apprise us of the structure of any experience given a sufficiently rich set of physical facts. Perhaps we might first need a qualitative bridging principle that tells us under what physical conditions consciousness is instantiated at all. Then we could apply our structural bridging principles to infer the structural properties of the experiences instantiated. However, developing such structural bridging principles requires us to overcome some big challenges. On the face of it, the structure of experience is radically different from the structure of the physical correlates of experience. It is plausible that there are systematic principles that connect physical structure to

phenomenal structure, but it is unclear exactly what these principles might look like. Developing such principles requires the empirical work of investigating the physical correlates of consciousness and evaluating different hypotheses, and the theoretical work of interpreting the data and developing generalized theories.²⁴ And even when we do acquire a plausible set of structural bridging principles for our own experiences, there is a question of whether we would be justified in extending those principles to the experiences of other creatures. As we venture away from the familiar cases, our evidential support weakens.

The most important challenge lies in attaining a better understanding of the structure of experience itself. Unless we have a grip on which structural properties experiences actually instantiate, we cannot develop structural bridging principles that link those structural properties to physical properties, or even appreciate which forms the objective phenomenal facts might take. To make progress, we need systematic tools for analyzing and modeling the structure of experience. In my view, this is one of the most important tasks facing the investigation of experience, and one of the most promising paths towards progress.

§ 5 | CONCLUSION

We started with a puzzle set introduced by Nagel. From there, we have developed a framework that enables us to examine the question of objective phenomenology with greater precision. The subjectivity of qualitative facts is explained by the subjectivity of pure phenomenal concepts, which itself is explained by the limits of introspection, extrapolation, and recombination. Structural facts, on the other hand, are objective because structural concepts and the concept EXPERIENCE are objective.

Structural facts are also epistemically interesting. Even though degree of generality correlates with degree of objectivity for qualitative facts, structural facts can be highly specific yet perfectly objective. Even though objective facts about the physical world are discoverable using solely third-person methods, discovering

²⁴ See Teller [1984] for an example of how this kind of project might begin.

structural facts about experience requires first-person methods because of the structural explanatory gap. And, of course, even though qualitative facts are typically subjective, structural facts are perfectly objective.

The principal goal of this paper has been to argue that an objective phenomenology is possible. To develop an objective phenomenology further, we need a better understanding of what kinds of structure experiences actually have and the principles connecting phenomenal structure to physical structure. As we make progress on those questions, we will better appreciate the limits of what we can understand and discover about the experiences of other creatures.

REFERENCES

- Ainsworth, Peter M. (2009). Newman's objection. *British Journal for the Philosophy of Science* 60 (1):135-171.
- Alter, Torin (2013). Social Externalism and the Knowledge Argument. *Mind* 122 (486):fzt072.
- Atkins, Richard Kenneth (2013). Toward an objective phenomenological vocabulary: how seeing a scarlet red is like hearing a trumpet's blare. *Phenomenology and the Cognitive Sciences* 12 (4):837-858.
- Ball, Derek (2009). There are no phenomenal concepts. *Mind* 118 (472):935-962.
- Block, Ned (2002). The harder problem of consciousness. *Journal of Philosophy* 99 (8):391-425.
- Block, Ned (2015). The Puzzle of Perceptual Precision. *Open Mind*.
- Carnap, Rudolf (1928). *The Logical Structure of the World*. Berkeley: University of California Press.
- Chalmers, David J. (2003 a). Consciousness and its place in nature. In Stephen P. Stich & Ted A. Warfield (eds.), *Blackwell Guide to the Philosophy of Mind*. Blackwell. pp. 102-142.
- Chalmers, David J. (2003 b). 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. (2004). How can we construct a science of consciousness? In Michael S. Gazzaniga (ed.), *The Cognitive Neurosciences Iii*. MIT Press 1111-1119.
- Chalmers, David (2017). The Combination Problem for Panpsychism. In Godehard Brüntrup & Ludwig Jaskolla (eds.), *Panpsychism*. Oxford University Press.

- Corwin, Thomas (2001). "Sensory adaptation: Tunable colour vision in a mantis shrimp". *Nature*. 411 (6837): 547–8.
- Cronin TW, Marshall NJ (1989). "A retina with at least ten spectral types of photoreceptors in a mantis shrimp". *Nature*. 339 (6220): 137–40.
- Fechner, G., 1860, *Elements of Psychophysics*, H.E. Adler (trans.), New York: Holt, Reinhart & Winston, 1966.
- Frigg, Roman & Votsis, Ioannis (2011). Everything you always wanted to know about structural realism but were afraid to ask. *European Journal for Philosophy of Science* 1 (2):227-276.
- Husserl, Edmund. 1913. *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy*. Trans. F. Kersten. The Hague: Nijhoff.
- Johnston, Mark (2007). Objective mind and the objectivity of our minds. *Philosophy and Phenomenological Research* 75 (2):233–268.
- Kolarik, Andrew J., Cirstea, Silvia, Pardhan, Shahina, & Moore, Brian C. J. (2014). A summary of research investigating echolocation abilities of blind and sighted humans.
- Levine, Joseph (1983). Materialism and qualia: The explanatory gap. *Pacific Philosophical Quarterly* 64 (October):354-61.
- Lockwood, Michael (1993). The grain problem. In Howard M. Robinson (ed.), *Objections to Physicalism*. Oxford University Press. pp. 271-291.
- Mensch, James R. (2000). An objective phenomenology: Husserl sees colors. *Journal of Philosophical Research* 25 (January):231-260.
- Nagel, Thomas (1974). What is it like to be a bat? *Philosophical Review* 83 (October):435-50.

- Nagel, Thomas (1986). *The View From Nowhere*. 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.
- Poincaré, H. ([1913] 1946). *The Value of Science*, translated by George B. Halsted. In H. Poincaré (Ed.), *The Foundations of science: Science and hypothesis, the value of science, and science and method*. Lancaster: The Science Press.
- Quine, W. (1968). I. Lakatos & A. Musgrave (Eds.), *Problems in the philosophy of science*. Amsterdam: North-Holland Publishing Company.
- Rabin, Gabriel (2011). Conceptual mastery and the knowledge argument. *Philosophical Studies* 154 (1):125-147.
- Sider, Theodore (2011). *Writing the Book of the World*. Oxford University Press.
- Teller, D. Y. (1984). "Linking propositions". *Vision Research*. 24 (10): 1233–1246.
- Watzl, Sebastian (2017). *Structuring Mind. The Nature of Attention and How it Shapes Consciousness*. Oxford, UK: Oxford University Press.