

CHAPTER 10
Materialism and the Epistemic
Significance of Consciousness

GEOFFREY LEE

Chapter Overview

Conscious experiences are commonly thought to have special epistemic properties: for example, they are thought to acquaint us with themselves and with external objects and events and enable us to form justified beliefs about ourselves and about the world. This chapter argues that if reductive materialism is true, this gives us reason to doubt that consciousness is unique in having these special epistemic properties. We tend to think that it *is* unique: for example, we tend to think that if we were turned into functional zombies by having our biological neurons replaced by artificial silicon components, this would put us in a massively impoverished epistemic state. I argue that this belief is grounded in a belief that there is a deep natural divide between beings that have consciousness and beings that lack it. But if reductive materialism is true, then it is wrong to think that there is any such deep natural divide, suggesting we should revise our belief in the unique epistemic significance of consciousness. I consider several strategies someone could adopt to resist this conditional conclusion. For example, they could hold that what it *is* to be conscious is to have states with this special significance. I argue that these strategies fail. If reductive materialism is true, then being a zombie can be just as good as having consciousness (at least epistemically).

Introduction

Conscious states are often thought to have great significance as a source of knowledge, both of the environment around the conscious subject and of the subject's own mind. Conscious perception is thought of as an interface

between belief states and the environment, presenting us with the local world in a way that allows us to form knowledgeable judgments about it. And conscious states themselves seem known to us in a peculiarly direct way—they are self-illuminating or self-revealing, the intimate epicenter of everything that is known to us. The special kinds of epistemic significance that attach to consciousness will be my subject here. More specifically, I am interested in whether conscious states are *unique* in the kind of epistemic role they play, or whether something other than consciousness could do the same work. Could a completely unconscious being have internal states with much the same epistemic significance as our conscious states?

We certainly have intuitions that suggest otherwise (Siewert, in [Chapter 9](#) of this volume, does a good job of making these vivid). For example, it seems plausible that a blindsighted subject who is able to reliably guess what is in the “blind” part of her visual field without consciously perceiving it, does not know about this part of the environment in anything like the way that conscious experience would allow. One explanation for this is that conscious experience is *necessary* for a certain kind of epistemic relationship to the environment. Despite such intuitions, I will argue that if materialism is true, we have reasons to doubt the idea that consciousness is unique in its epistemic properties. On the materialist view consciousness is a high-level physical or functional feature of the brain (or some larger system). A being with a quite different cognitive architecture might lack this (seemingly) special property but have internal states that play a similar role, including in the provision of reliable mental representations of the environment and their own mind. I will argue that despite lacking conscious awareness, such a being might be in an epistemic situation just as good as ours.

By contrast, Siewert argues that consciousness is unique in its epistemic significance, grounding the justification of perceptual and introspective beliefs, and our understanding of perceptual and introspective concepts. He claims that a philosophical zombie could *not* have internal states that play the same epistemic role, and even goes as far as to claim that consciousness is necessary for having any mental states at all. He rejects reductive materialism, so perhaps his overall view is not incompatible with my conclusions here. However, the arguments he offers do not seem to depend on rejecting reductive materialism: so if I am correct, these arguments are inadequate, even if his view remains a live possibility.

I will discuss the following three kinds of epistemic significance that seem to attach to conscious states: *doxastic significance*, *intrinsic epistemic significance*, and *concept-grounding significance*. *Doxastic significance* has to do with the ways conscious states are able to confer justification (or other positive epistemic status) on beliefs. For example, it is often held that conscious perceptual states can contribute to justifying perceptual beliefs, as when an experience of the color of a surface helps to justify you in judging what color it has. *Intrinsic*

epistemic significance has to do with the ways in which conscious states are in themselves epistemically valuable. The idea is that conscious states intrinsically involve a kind of epistemic relation either to the environment or your own mind, independently of what you believe. Some theorists elaborate in terms of the idea of *acquaintance*—an experiential relation that directly confronts you with something and enables you to know it in a special direct way. Intuitively, even if you lacked the capacity for belief, conscious acquaintance would still provide a form of epistemic contact with the world. Finally, *concept-grounding significance* has to do with the role of conscious states in enabling a subject to grasp concepts that refer to internal or external items, such as the objects or events one is acquainted with in having a perceptual experience. For example, a conscious perceptual experience of the color of a surface may be part of what enables you to think of the color as “that color.” I will focus particularly on doxastic significance, arguing that if materialism is true, consciousness is not unique in its ability to have this kind of significance: there is a family of states that an unconscious being could enjoy, which have the same doxastic significance as if they were conscious. Intrinsic and concept-grounding significance will play a role though, because it might be claimed by an opponent that they explain why consciousness is unique in its doxastic significance. I will return to this idea subsequently.

Epistemic significance is not the only kind of significance thought to attach to consciousness, and it is worth briefly discussing the other varieties before proceeding. One of them is moral or practical significance. If we believe that certain entities have intrinsic value, one plausible candidate would be experiences like pleasures and pains (and even if we don’t believe in intrinsic value, it is certainly true that we *care* about pleasure and pain). More strongly, it might be held that conscious experience is necessary part of the ground for *all* value in the world—that without conscious experience nothing would matter. Despite its obvious importance, this kind of significance won’t be my focus here (although I suspect many of the points here about epistemic normativity have analogues in the practical case).

A third kind of significance, whose connections to the epistemic significance of consciousness will be explored here, is what might be called “natural significance.” Consciousness (or specific kinds of consciousness) is naturally significant if it is a highly “natural” property, or if there is a deep natural divide between the beings that have consciousness and those that do not (this being one reason why consciousness might be thought to have special scientific interest). Following Armstrong and Lewis (see Armstrong, 1978; Lewis, 1983), many philosophers, including myself, believe that properties are distinguished by some objective measure of naturalness, and that one goal of science is to articulate what these natural “joints” in the world are. For example, we want to say that two electrons are objectively more similar than an electron and a proton, even though there are ways of understanding what a “property” is on

which all particles share equally many properties (e.g., they are equally many sets of possible objects to which they belong). It seems that some of these “thin” properties count more toward objective resemblance than others, and this gives an important role for the concept of “naturalness.”¹ Being conscious seems to be a particularly striking objective resemblance between objects, and for this reason, it might be supposed that consciousness is a highly natural property.

Although the natural significance of consciousness won’t be my main focus in this chapter (it is the focus of Lee [forthcoming]), it will play an important role in the argument; in particular, my discussion will be organized around an argument to the effect that *if* consciousness lacks special natural significance, then it lacks special epistemic significance (in certain sense).

I am going to challenge the epistemic significance of consciousness, but it is important to be clear on the strength of the challenge I am offering. For any kind of significance that allegedly attaches to consciousness (epistemic, practical, or natural), I would distinguish three grades of challenge to it. First, one could deny that *anything* is significant in the relevant respect. For example, one might be skeptical about whether there is any objective distinction of naturalness between properties (e.g., Goodman, 1983; Taylor, 1993) or be a nihilist, denying that anything is intrinsically valuable, or be skeptical about whether there is any objective sense in which some cognitive states are epistemically significant—(e.g., more “rational,” “justified,” or epistemically valuable) than others (e.g., Field, 2009). Second, one could think that although some things are significant in the relevant respect, consciousness does not have the relevant significance, or at any rate is *less* significant in the relevant respect than other entities—for example, one might think that only beliefs can have epistemic value, and conscious experiences only have derivative epistemic value, for example, as their reliable causes (Davidson, 1986). These are not the challenges I will be primarily concerned with here. Instead, I am concerned to deny the idea that consciousness has *unique* significance, in the sense that there are no similar properties—properties that an unconscious being could have that are *equally* significant (in some given respect perhaps). Consciousness does not “stand out from the crowd.”

To get a sense of this view, imagine you are faced with the following dilemma (a similar case is given by Dainton [2008, p. 181]): You are suffering from a degenerative neural condition. If it is left untreated, you can expect only another year of relatively normal cognitive functioning, and you will be dead in two years. Scientists offer you a treatment that will not only extend your lifespan but also enhance most of your cognitive processing (your ability to access stored information, reason with it, etc.) while leaving most of your current mental states (your personality, beliefs, desires, intentions, etc.) intact. Your friends will meet a brighter, sharper, version of your former self. There is one large catch, however—the treatment involves gradually replacing your

neural circuitry with a functionally equivalent silicon proxy, the result being that you will completely lose your capacity for conscious experience (it is not specified how it is *known* that this will be the result—we just stipulate that it is known). You will still “perceive” the world in some sense, but your perception will completely lack the characteristic subjective feel that makes mental states “conscious” in the sense that we seem to care about (often referred to by philosophers as “phenomenal consciousness”).

Should you accept the treatment? I think most people would not accept it. Having conscious experience seems to be necessary for having the kind of life we care about, and being without it—having the “light inside” permanently extinguished—seems pretty much as bad as ordinary death. Notice, however, that this reasoning is only sound if the silicon zombie states aren’t as *good* as conscious states in some sense. Suppose, however, that the zombie states have the same moral, epistemic, and natural significance as our conscious states do. Then it is hard to see how they are less desirable, and it looks like we ought to accept the operation. My claim will be that if materialism is true, then there are zombie states that have the same natural and epistemic significance as conscious states, and to that extent at least, are just as good. I call this view “deflationary pluralism.”

Some terminology is helpful in order to define the view more precisely. Let us say that a state is “consciousness-like” if the functional constraints it puts on the cognitive architecture of an organism at least superficially resemble those associated with consciousness. We can say that a being is *pseudo-conscious* if they are in an internal state that is consciousness-like but are not conscious. More specifically, we can distinguish between internal states that are *folk-consciousness-like* and *empirically consciousness-like*.² A being is *folk-pseudo-conscious* if she is in an internal state that satisfies the functional criteria that ordinary people associate with consciousness (e.g., the ability to report the relevant states), but is not conscious. A being is *empirically pseudo-conscious* if she is in an internal state that satisfies many of the architectural details that have been found empirically to be associated with consciousness (e.g., having perceptual modules feeding into a central pool of information used to reason and initiate high-level motor plans) but is not conscious. We can make the same distinctions for determinate kinds of conscious states. For example, we can talk about pseudo-pain, folk-pseudo-pain, empirical pseudo-pain, and so on.

Let us further distinguish *quasi-consciousness* and pseudo-consciousness. A being is quasi-conscious just if she is pseudo-conscious and her form of pseudo-consciousness is equally significant as consciousness. We may want to distinguish different kinds of significance here: we can talk about epistemic quasi-consciousness and practical quasi-consciousness, and we could make more fine-grained distinctions. The distinctive claim of the deflationary pluralist is that, as well as consciousness, there is a whole family of ways of being quasi-conscious. The claim I am interested in here is that *epistemic*

quasi-consciousness is possible: an unconscious being that superficially duplicated us could be equally justified in her beliefs and equally well “acquainted” with her environment as us.

The deflationist’s claim that forms of pseudo-consciousness can have the same significance (natural, epistemic, or practical) as real conscious experience can be made vivid by considering disputes about problematic cases such as lobsters or aliens, where it is unclear whether a being has conscious experience. The deflationist may regard these as empty questions, in the sense that even if these being aren’t conscious, they may have a consciousness-like internal state that is equally significant as consciousness, and so the dispute about whether they are conscious does not concern where some special deep boundary between us and other creatures lies. Another way of looking at it is that even if our concept of “consciousness” does not include lobsters, perhaps a variant on our concept that did include them would be no worse than our concept, in the sense that it would not miss a deep boundary that our concept *does* mark. (See Sider [2011] for a development of the notion of an empty question along exactly these lines.)

Deflationary pluralism implies, and so depends on, the weaker claim that a being might superficially duplicate us without being conscious. Some theorists, *superficial functionalists* (Block, 2002), might deny this, holding that having a functional organization superficially like that of a conscious human is sufficient for being conscious. It is worth noting, however, that a quite radical version of superficial functionalism is required to imply that no form of pseudo-consciousness is possible. The problem is that there is more than one way to superficially functionally resemble a conscious system, and they could only *all* be sufficient for consciousness if consciousness is an extremely abstract property that is determined by every such mode of resemblance (or it is highly disjunctive). For example, a human baby and an android might both have consciousness-like internal states but resemble conscious beings in very different ways, and as a result, they have very little in common with each other (see McLaughlin, 2003). For this reason, I believe superficial functionalism—or at least the radical kind that would undermine deflationary pluralism—to be an implausible view, which I will set aside here, acknowledging that it warrants more discussion. I will assume that at least *some* forms of pseudo-consciousness are possible.

The Grounding Argument for Deflationary Pluralism

I will organize my discussion around an argument that links together the natural significance and epistemic significance of consciousness (or normative forms of significance more generally, including practical and epistemic significance). I call it the natural-normative grounding argument, or the “grounding argument” for short. Let us say that consciousness has “strong epistemic significance” if epistemic quasi-consciousness is impossible (and let us define “strong natural significance” similarly). The argument looks like this:

- (1) Consciousness has strong epistemic significance only if it has strong natural significance.
- (2) If reductive materialism is true, consciousness does not have strong natural significance.
- (3) Conclusion: if reductive materialism is true, then consciousness does not have strong epistemic significance.

To get a feel for the argument, it is helpful to first consider the analogous argument for the moral or practical significance of consciousness. Consider conscious pain compared with a zombie analogue of conscious pain—an unconscious state with a superficially similar functional role. Intuitively, there is a massive difference between these two things. One is a mere functional simulacrum of a kind of state that is very special in virtue of having “phenomenality” or “subjective feel.” Moreover, although conscious pain would be judged by many to be an inherently bad thing, many would think that the zombie analogue is not at all bad in the same way. What is the relation between this apparent normative fact and the apparent fact that there is a “deep” difference between pain and zombie pain? On one view (an epistemic version of which is discussed in detail subsequently), the deep difference just *consists* in the fact that pain is bad and zombie pain isn’t. However, I don’t think this is the intuitive view—intuitively, pain is bad and zombie pain isn’t in part *because* they are massively different kinds of things. This massive difference can’t therefore *consist* in the normative difference, but must rather be something else that (at least partly) *grounds* it. The obvious candidate is that there is a deep *natural* difference between pain and zombie pain.

But *is* there a deep natural difference between pain and zombie pain? Arguably, if reductive materialism is true, then there is not.

Reductive materialism (as I understand it) is the view that consciousness is identical to a high-level functional or physical property of the brain or some larger system (ditto for more determinate kinds of consciousness). Thus, consciousness and its determinates have “real definitions” (specifications of what it is to have the property in terms of these other properties) in terms of more basic physical and functional properties, implying that facts about consciousness hold wholly in virtue of facts about the physical and functional properties of conscious systems.

If reductive materialism is true, it might be argued that real pain won’t be in any way special from an objective perspective compared with pseudo-pain—they are both fairly similar high-level physical-functional properties of a psychological system. To make this idea vivid, imagine a Martian scientist comparing the physical-functional state that *is* pain to the analogue zombie state in, say, a silicon-based life-form. She will see very abstract high-level similarities in the mental architectures of the two life-forms, with completely different realizations of these architectures. It is implausible that the Martian

scientist will think that any great natural significance attaches to one realization over the other—it is not as if the realization that enables consciousness will be seen as having a magical glow. In this sense, pseudo-pain might be rather *similar* to real pain. If this is right, we may on reflection want to retract our initially confident judgment that pseudo-pain can't be just as *bad* as real pain, which seems to be partly based on the assumption that they are *not at all* similar.

Compare this with the situation if reductive materialism is false. The alternatives are property dualism and “non-reductive materialism,” views on which consciousness and its determinates are primitive irreducible properties whose instantiation does not depend in any intelligible way on lower-level facts: they are metaphysical brutes (I think it is fair to characterize these views this way, even if the primitivist thinks that mental properties supervene on lower-level properties—as on a “non-reductive materialist” view—since this will be brute supervenience of one kind of primitive property on other primitive properties.) On a primitivist view like this, the Martian scientist's knowledge of the physical and functional differences and similarities between us and pseudo-conscious aliens won't reveal the fact that we are distinguished in a deep way from the aliens by enjoying primitive phenomenal properties.

In Lee (forthcoming), I unpack the Martian scientist metaphor in detail, discussing the different interpretations of “natural significance” one could have and the inference from “the Martian won't see a big difference in natural significance between consciousness and pseudo-consciousness” to “there *is* no big difference” (assuming reductive materialism). A thought experiment that helps one see the point involves imagining a spectrum of cases linking a conscious being to a pseudo-conscious being like a silicon android. Each being in the spectrum has the same high-level architecture—they are in the same “consciousness-like” state—but there is a small difference in realization from one case to the next (e.g., a few neurons might be replaced by functionally equivalent silicon components). Even though there is consciousness at one end of the spectrum, and no consciousness at the other end, the shift from one case to the next will seem trivial from an objective perspective—the perspective of the Martian scientist. So there appears to be no point at which we reach a monumental divide between cases of consciousness and cases of mere pseudo-consciousness, supporting the view that consciousness and pseudo-consciousness have equal natural significance. (Note: I am intending this thought experiment as more of an aid to intuition than as an argument for deflationism.)

It is natural to object that even if the Martian scientist sees no big divide between us and unconscious silicon beings, we can appreciate that there *is* a big divide given the understanding of consciousness we get “from the inside.” In brief, my response to this is that if reductive materialism is true, then the natural significance of consciousness is not best appreciated “from the inside” but rather is optimally assessed given knowledge of what consciousness really

is—that is, given knowledge of its real definition in physical/functional terms.³ Given reductive materialism, the objective, not the subjective, perspective is privileged for assessing the natural significance of consciousness—that is, the Martian scientist is in a better position to understand the natural significance of consciousness than we are.

For the purpose of this discussion then, I will assume that premise 2 is correct—that if reductive materialism is true, then there is no special natural significance that attaches to consciousness but not to any form of pseudo-consciousness. I am interested in what this assumption implies about the other forms of significance alleged to attach to conscious states, in particular epistemic significance.

If the grounding argument works for the practical significance of consciousness, then it is quite plausible that it works in the epistemic case as well. If you think that a conscious percept can acquaint you with the environment, enable you to think about features of the environment and form judgments about them, whereas an unconscious state can do none of these things, intuitively this difference is *grounded* in the fact that an unconscious state is a quite different type of thing from a conscious one. It is not what the difference *consists in*. I think this grounding argument may be more controversial in the epistemic case than in the moral case, but nonetheless it has force. Most of what follows will be concerned with discussing how viable it is to reject it (i.e., to reject premise 1), holding that consciousness has special epistemic significance that is *not* grounded in special natural significance.

I will now proceed by considering two different kinds views, each of which, if correct, would undermine the grounding argument. First, one could take the view that to be conscious just *is* to be in a state with a certain kind of epistemic significance; so it is in the very nature of consciousness that epistemic quasi-consciousness is impossible. I call this view *normativism*. The alternative to a normativist view is to hold that although consciousness is not defined in terms of its epistemic significance (which I interpret as implying that it is a natural property of the kind that features in causal laws and explanations of the kind found in science), it nonetheless has strong epistemic significance—for example, the norms governing belief formation in fact imply that conscious experience is required to have rational belief. I call this view *strong epistemic naturalism*. I think any attempt to defend the strong epistemic significance of consciousness will be in either one of these categories: I now consider them in turn.

First Response: Normativism—Being Conscious *Just Is* Having States with a Certain Epistemic Significance

According to the grounding argument, if consciousness has strong epistemic significance, this is partly grounded in the fact that there is a big natural divide between the conscious beings and the rest (premise 1). One way around this is

to claim that consciousness is just *defined* in terms of its epistemic significance, in such a way that having states that function like conscious states, and that have the epistemic role of conscious states, is what *makes* you conscious. On this view, the big divide between conscious and unconscious beings is not a natural one that ought to be salient to a Martian scientist, but rather one defined in epistemically normative terms; it is the divide between having states with a certain kind of (normative) epistemic significance and lacking such states. As well as conflicting with the natural/normative grounding claim (premise 1), this view can be understood as directly incompatible with the existence of epistemic quasi-consciousness (i.e., the conclusion of the argument), implying that the grounding argument must be unsound. I think this rejection of grounding is counterintuitive, but the view is still worth discussing in some detail.

I will assume that the kind of epistemic significance that is supposed to define consciousness is *doxastic*; that is, it has to do with the way conscious states can contribute to justifying beliefs, such as beliefs about your environment, or beliefs about your own mind. We can think of this in terms of the existence of epistemic *norms* governing the formation of belief. Belief formation involves a transition from a prior to a posterior mental state, where the mental state involves a set of beliefs and other relevant states, such as conscious experiences. A necessary condition on a belief being justified will be that it is formed in accordance with the relevant transition norms. The normativist thinks that the conscious states a subject has are, at the very least, *relevant* to what transitions they are justified in making; that is, they are relevant to whether the subject is following the relevant transition norms.⁴ For example, if I have a conscious perceptual experience as of a yellow feather, this may contribute to my being justified in forming the belief that there is a yellow feather before me, in accordance with a norm governing the formation of perceptual belief.

We can call this kind of relevance to justification *weak doxastic significance*. Consciousness having weak doxastic significance (as I understand this) is consistent with forms of pseudo-consciousness also having it. There might be alternative norms that a zombie could follow, norms that ground the rationality of transitions from pseudo-conscious states to beliefs.

We can say that consciousness has *strong doxastic significance* if there are no such analogues of consciousness-involving norms—that is, if the correct epistemic norms are *consciousness-requiring*. Normativism is supposed to imply this stronger significance because it is the view that following the norms that govern certain kinds of psychological transitions (e.g., those from perception-like states to beliefs) is *sufficient* for being conscious. We can understand this as the view that consciousness is defined by a “functional-normative” role:

Normative Functionalism: Being conscious (or having some specific type of consciousness) is the property of being in a state that plays a certain functional-normative role.

A functional-normative role is like a regular causal role, except that the mental transitions that specify the functional role are taken to have a normative aspect—they are (by definition) transitions that one *ought* to perform if one is in the relevant mental state. For example, a normativist about perceptual consciousness might define it in terms of a norm telling one to “endorse” the content of a perceptual experience one is in; for the normativist, having a conscious perceptual experience might be taken to partly *consist* in being in some state that is governed by a content-endorsement norm.

Playing a normative-functional role is a stronger requirement than playing the corresponding causal-functional role, precisely because it adds in a normative requirement. Satisfying the functional requirement might not be enough to satisfy the normative requirement: for example, it might be that to be conscious—and thereby to have states capable of grounding justified beliefs about yourself and the environment—the relevant functional property has to be realized on a certain kind of neural hardware (this is one reason why normative functionalism is compatible with the possibility of pseudo-consciousness). On the other hand, satisfying the functional requirement might be sufficient for consciousness; the view would still differ from regular functionalism in the sense that it holds that the nature of consciousness is partly normative, having to do with its role justifying belief.

What exactly are the epistemic norms that govern consciousness (and on this view define it)? It will be helpful at this point to distinguish two kinds of views of the normative role of conscious experience that a believer in strong normative significance might have. First, they might believe the following:

Strong Foundational Significance Thesis: Necessarily, if a belief is justified non-inferentially (i.e., not on the basis of other beliefs), part of the ground for this justification is the existence of a conscious state.

The picture here is one of a grand web of beliefs, most justified inferentially through a chain of other beliefs, but some providing a foundation on the periphery, where justification by inference from other beliefs gives out. These non-inferential beliefs are justified by transitions from conscious experiences (e.g., conscious perceptions), which thereby ground the whole structure. We can understand the grounding claim here as implying that if you lack conscious states, then your non-inferential beliefs are not justified: *only* consciousness can provide such a foundation. Thus a zombie’s non-inferential beliefs, and any other beliefs whose justification depends on that of non-inferential beliefs, are not justified. This suggests that a zombie can’t have *any* justified beliefs. Having said this, it might be plausible to restrict the view, so that only a certain class of non-inferential beliefs, such as introspective psychological judgments and non-inferential empirical judgments, are held to require conscious experiences to be justified. This would allow for the fact that some non-inferential

beliefs, such as a priori beliefs about mathematics, do not seem to be justified by experience. On this restricted version, a zombie might be able to form justified beliefs about mathematics and other a priori matters, but would not be able to have justified beliefs about its environment or its own psychological states.

Phenomenal consciousness might play a more general justificatory role than injecting from the foundations, as it were. For example, *access internalists* hold that the fact (or proposition) that p can only be part of a subject's reason for believing that q if the subject has "access" to p in some sense. For example, Freudian repressed beliefs do not seem to be rational grounds for other beliefs—if I repressedly believe that Britain is not the world's greatest nation, I can't use this as a premise in reasoning. One way of understanding the relevant notion of "access" is in terms of a subject's phenomenal states:

Phenomenal Access Internalism: Necessarily, the fact that p can only be part of a subject's reason for believing that q if p is available to phenomenal awareness (a similar view is defended by Smithies [2012a]).

We can say that p is "available to phenomenal awareness" just if p is either the content of a belief that is consciously accessible, or p is the content of a perceptual experience, or is an introspectible fact about a phenomenal state (such as the fact that it is an intense pain, or a visual experience of blue). Phenomenal access internalism implies that consciousness has a more general role conferring justification than merely injecting justification at the foundations because it requires that *anything* that is part of a subject's reasons for believing something is consciously accessible, including the contents of beliefs that have been formed by inference from other beliefs.

These views combine with normativism as follows. For the foundationalist normativist, being conscious just *is* being capable of mental transitions that satisfy the relevant foundational norms. For the phenomenal access internalist normativist, being conscious will be the property of having states that are subject to the relevant access internalist norms⁵ (e.g., they are available for justified self-ascription).

Is there anything that can be said to independently justify the normativist view? By way of analogy, consider the view of philosophers like Davidson (1973), McDowell (1985), and Dennett (1971), who think there is a normative aspect to the ascription of propositional attitude states. The idea is that to treat a subject as having intentional states like beliefs and desires is to treat them as rational agent with states governed by certain epistemic norms; for this reason, intentional ascription is subject to a "principle of charity," according to which subjects of belief and desire are, *ceteris paribus*, rational. One can understand this as the view that belief and desire states are at least partly *defined* in terms of their rational relations to each other, and so belief/desire ascriptions are at least

partly normative statements (although there are other interpretations of the view⁶). One can imagine a theorist who thinks that a similar point is correct for conscious experience—that we can only make sense of it as a content-bearing state by conceiving of it as standing in certain rational relations to other mental states, rational relations that partly define what it is to have the conscious state.

In response to this, there seems to be an important disanalogy between conscious experiences and propositional attitudes. Our understanding of what consciousness is clearly does *not* depend on a holistic understanding of how it fits into a whole cognitive system, including how it is rationally related to other kinds of mental states. In particular, we know what conscious states *are* by simply *introspecting* them, a process that doesn't require any background theoretical understanding of other aspects of the mind and their rational relations to consciousness. This seems an obvious contrast with our understanding of what beliefs are, which may require a grasp of rational belief-desire psychology (even if we can introspect beliefs). If it is true that we do not conceptualize conscious states in terms of their rational relations, it would be surprising if they are in fact individuated this way; certainly, the burden of proof is on the normativist to show this.

As well as lacking clear independent support, there are direct arguments against normativism about consciousness. In particular, there are problems that arise when we look at how it plays out on different views of the metaphysics of normative facts. I will argue that however we develop the view, it is either incompatible with reductive materialism or otherwise problematic.

Normativists who are realists about norms in general, either think that there are irreducible normative facts or that normative facts can be reduced to non-normative facts somehow (e.g., they might say that for a norm to hold is for people to *accept* that it holds).⁷ Normativists who believe that there are irreducible normative facts will construe facts about consciousness as holding partly in virtue of these irreducible normative facts. But this is incompatible with reductive materialism: I am thinking of reductive materialism as implying that facts about consciousness hold wholly in virtue of facts about the functional and material properties of conscious systems, and therefore not in virtue of irreducible normative facts.^{8,9}

If this is right, then for realist normativism to be consistent with reductive materialism, the normativist will have to be a reductionist about normative facts. However, normative reductionism is independently very implausible: for example, if we say that normative facts are a special kind of natural fact, such as the fact that certain people *accept* certain norms—it is hard to see how they are genuinely normative, or have normative force.¹⁰ Furthermore, whatever natural fact about us constitutes consciousness-requiring norms obtaining, it seem likely that a similar natural fact will obtain for pseudo-conscious beings; for example, even if *we* accept that consciousness-requiring norms obtain, it will also be true that *they* accept that pseudo-consciousness-requiring norms

obtain. The realist reductionist normativist will therefore find it hard to maintain that there is any deep asymmetry between us and pseudo-conscious beings.

What about the combination of normativism with *anti*-realism about norms? Normative anti-realists generally either hold an error theory about normative claims (holding the view that all normative claims are false) or are expressivists (holding the view that normative claims are not really genuinely fact stating).¹¹ Thus the error-theoretic normativist will say that statements about consciousness are always false, because they involve normative commitments that nothing ever satisfies (because there are no normative facts), whereas expressivist normativists will say that statements about consciousness are not properly fact stating, but rather serve to express non-cognitive attitudes like desires. Are these views at all plausible? A consciousness error theory seems obviously false, and the expressivist view, although interesting and worthy of more discussion, is quite counterintuitive. It is plausible that when we talk about conscious experience we are intending to represent reality as being a certain way, and not merely emote or perform some other kind of non-descriptive act.¹² In this context however, a more fundamental objection to both anti-realist views is simply that by denying consciousness any kind of robust existence, they are also denying it any kind of robust significance.

To sum up, we can think of normativism as the view that conscious states are by definition those that play a certain role justifying beliefs; this will likely either be a foundational role justifying non-inferential beliefs or a role in providing the propositions about world to which the subject has the kind of access needed for them to provide reasons for belief. Despite being an obvious way to defend the strong epistemic significance of consciousness, we saw that it is hard to give the view independent motivation—for example, the considerations that support a normative view of propositional attitudes don't translate well to conscious states. We also saw that different views of the nature of normative facts lead to different version of normativism that each are problematic in one way or other. Let us therefore move to consider a different response to the grounding argument.

Second Response: Strong Epistemic Naturalism—Consciousness Is a Relatively Shallow Natural Property, but Nonetheless Unique in Its Epistemic Significance

The other kind of view that might allow one to defend the strong epistemic significance of consciousness by rejecting the grounding claim (premise 1) is what I'll call strong epistemic naturalism (SE naturalism). On this view, consciousness is *not* defined in normative terms—it is a completely natural phenomenon, such as a complex functional property—but nonetheless it has strong epistemic significance. Because the SE naturalist rejects the grounding claim, she denies that this epistemic significance is grounded in a deep

natural divide between conscious and unconscious beings (consciousness is not an outstandingly deep natural property). Rather, the fact that the correct epistemic norms are consciousness requiring is either brute or grounded in something other than the nature of consciousness itself. The view therefore cannot explain the strong epistemic significance of consciousness in the same way that normativism can, but it has the advantage of being able to treat consciousness as a completely natural phenomenon.

I will assume here that a zombie can have propositional attitudes and other mental states (e.g., pseudo-conscious perceptual states) with contents at least analogous to the contents of our mental states (I say “at least analogous” because their beliefs about pseudo-consciousness will not have exactly the same content as our beliefs about consciousness). Siewert, following Searle (1989), rejects this assumption, although his case for it depends on a prior rejection of the epistemic significance of any zombie states. A full argument for my conditional conclusion would consider in detail the claim that consciousness is necessary for intentionality in general. Unfortunately, I do not have space to discuss this here: unsurprisingly, my view is that if there is no deep natural asymmetry between us and functional zombies (as I think is true on a materialist view), then they can at least have something extremely similar to our intentional states.

I will focus on a version of SE naturalism that involves the doxastic significance of consciousness, but the idea that it has intrinsic epistemic significance (e.g., that conscious states acquaint their subject with the environment or their own mind) or concept-grounding significance (conscious states enable subjects to form certain concepts, such as concepts of objects and properties that the experiences acquaint their subjects with) will be part of the discussion as well.

The version of SE naturalism that I’m interested in therefore involves the claim that there exists a set of norms governing the correct formation of belief, which together imply that conscious experience is *required* for certain kinds of beliefs to be justified. I’ll call this the “consciousness-loving” view. As before, I will assume that the consciousness lover is either a phenomenal foundationalist or a phenomenal access internalist. I will assume that their zombie-loving opponent believes in the existence of similar norms but thinks that these norms are consciousness-neutral, so that a pseudo-conscious zombie could also follow them. For example, instead of thinking that only phenomenal states can play a foundational role, they might hold that a class of pseudo-phenomenal states can also do this.

How might the consciousness lover justify their normative view? There appear to be three possibilities: (1) To hold that consciousness-requiring norms have a ground in certain non-normative facts, such as the fact that conscious states provide “acquaintance” with themselves or with external events (the “external grounding” strategy); (2) to hold that consciousness-requiring norms

are explained in terms of other epistemic norms, norms that do not themselves mention consciousness (the “normative grounding” strategy); or (3) to hold that certain consciousness-requiring norms obtain, but that this fact can’t be further explained—it is one of the consciousness lover’s basic commitments. I’ll call this the “no grounding strategy.”

Let me consider these three defensive strategies in turn, starting with the external grounding strategy. I can think of the three ways that consciousness-requiring norms might be thought to be externally grounded in non-normative facts: (1a) they are grounded in the intrinsic epistemic significance of consciousness, (1b) they are grounded in concept-grounding significance of consciousness, or (1c) they are grounded in what I will call the “introspective significance” of consciousness (a view defended by Smithies [2012a, 2012b]).

Again, to say that conscious states are “intrinsically epistemically significant” is to say that they are an epistemic good in themselves, independently of their role in conferring justification on beliefs. A likely elaboration of this idea would be that conscious states provide conscious *acquaintance* with objects and events in the environment or with the subject’s own mind, where acquaintance is an epistemic relation that can be enjoyed even by creatures incapable of belief. It might be thought that it is because zombies cannot be acquainted with the world or their own minds in the way that we can, that they cannot have knowledge or justified belief about these things in the way that we can. When it comes to grounding the norms governing belief, it might be said that it is a basic goal of belief formation to form beliefs based on evidence that is phenomenally accessible, and this basic norm is grounded in the intrinsic epistemic significance of conscious acquaintance. A comparison with moral norms might be illuminating. Why should we try to prevent pain? On one view, this is because pain is intrinsically bad, and this intrinsic badness explains why a pain-prevention norm obtains. Similarly, consciousness might have intrinsic epistemic goodness in a way that explains why certain consciousness-requiring epistemic norms obtain.

The zombie lover should counter-respond as follows: it is hard to see what this alleged “intrinsic” epistemic significance of conscious acquaintance could amount to. We are assuming that if reductive materialism is true, then conscious acquaintance does not have special *natural* significance; considered as an objective natural phenomenon, there is nothing special about conscious acquaintance over a pseudo-conscious functional analogue of acquaintance that could be enjoyed by a zombie. But the only other kind of significance we have a fairly clear understanding of is *normative* significance; conscious acquaintance might be significant because it is involved in the norms governing belief formation or epistemic enquiry more generally. But then this is not a separate, more “intrinsic,” kind of significance and therefore cannot plausibly *ground* or *explain* the correctness of the relevant norms. Compare with the

moral case again: one might object to the story told in the last paragraph by pointing out that the “badness” of pain can only be understood in terms of the fact that we ought to prevent pain; it therefore cannot ground this normative fact.

A second external-grounding approach is try to ground consciousness-requiring norms by appealing to the role of consciousness in the possession of certain concepts. Here I have in mind an approach like the one elaborated by Peacocke (1992, 2004), which attempts to explain the correctness of certain epistemic norms in terms of the conditions for possessing concepts, in this case “observational,” “perceptual,” or “recognitional” concepts, concepts whose mastery requires applying them to conscious experience in various ways. For example, it might be said that possessing the concept “red” requires being disposed to use this concept in judgments based on conscious experiences of red surfaces. According to this view, the correctness of certain consciousness-requiring norms for judgments involving “red” flows from the fact that these are the norms one must follow in order to *have* the concept in the first place.

Can this approach help the consciousness lover? One might think that there is nothing to stop a zombie from possessing observational concepts analogous to ours, but whose possession involves responding to pseudo-conscious states rather than conscious states, and thus following norms that are not consciousness requiring. Prima facie then, the concept-possession approach actually supports rather than hinders the zombie lover’s view. However, it might be replied that not all concepts are on a par—perhaps some concepts have possession conditions that make them in some way defective (e.g., Peacocke [1992] argues that this is the case with Prior’s [1967] concept “tonk”). Perhaps zombie observational concepts are *epistemically defective* because the zombie does not know what he is thinking about in the intimate way that conscious experience allows us to.

The problem with this is that we will probably need to appeal to an independent intuition about the epistemic role of consciousness to argue that only consciousness can ground non-defective observational concepts. This is another place where the intrinsic significance of consciousness might seem to play a role. For example, one might be attracted to the idea that conscious perception acquaints you with the environment, and the function of observational judgments is to exploit this acquaintance at the level of thought. On one elaboration (see, e.g., Campbell, 2002), this amounts to the acquaintance providing you with knowledge of the semantic values of your observational concepts, knowledge that is, intuitively, required for using concepts with genuine understanding of what they refer to. For example, when I think “this flower is beautiful,” I might know which flower I’m thinking about because I am enjoying a conscious visual experience of the flower. A zombie cannot think about the world in a way that is grounded in conscious acquaintance and so might lack this special kind of knowledge of what their thoughts are referring to. In

short, it might be that the doxastic significance of consciousness is explained by its concept-grounding significance, which is in turn explained by the intrinsic epistemic significance of conscious acquaintance.

This assumes again, however, that there is something intrinsically epistemically special about conscious acquaintance, as opposed to the pseudo-conscious acquaintance of a functional zombie—and the problem is that it is hard to understand what this means, if it is not a matter of natural or normative significance. This suggests that the concept-possession approach cannot on its own help the consciousness lover, although it does link doxastic significance with other forms of significance in an interesting way.

A third approach that is worth briefly mentioning is due to Smithies (2012a, 2012b). Smithies thinks that only introspectively accessible mental states can provide justification for our beliefs, a kind of access internalist view. Furthermore, he thinks that only phenomenally conscious states are introspectively accessible. Thus, he thinks the existence of consciousness-requiring norms can be explained in terms of an access internalist conception of justification (which he further elaborates in terms of the idea that a justified belief is one that would survive ideal critical reflection) and what we might call the “introspective significance” of phenomenal experience.

The problem with this approach is that our functional zombies do appear to have something like introspective access to their mental states, even though they lack consciousness. Smithies needs a reason why zombie introspection isn't as good as phenomenally grounded introspection. Giving such a reason will require, in effect, adopting one of the other strategies on the list since it involves explaining why the norms governing introspective belief formation are consciousness-requiring. In fact, his view seems to really be a kind of no-grounding view (a type 3 strategy): there are consciousness-requiring norms governing introspection that can't be further explained (he motivates them with blindsight-type examples but does not attempt to ground them). So his approach, although adding some flesh to a consciousness-loving picture, does not constitute a separate approach to the others I will consider.

Having looked at the external-grounding strategy, let's now briefly consider the normative-grounding strategy. Can consciousness-requiring norms be explained in terms of other, more basic norms? It is controversial what basic norms might govern belief formation, but plausible suggestions are external norms telling us to form true beliefs, or have beliefs formed by norms that reliably produce true beliefs, or more “internalist” norms, such as a norm telling us to have coherent beliefs. It doesn't matter for my purposes exactly what norms of this kind, if any, hold, because there is a general problem with this strategy. The problem is that a functional zombie can have beliefs that exactly mirror our beliefs, standing in the same causal and inferential relations, and formed on the basis of equally reliable processes linking them with the world outside the head. So, apart from norms explicitly mentioning consciousness, it seems

likely that they can satisfy *any* plausible general epistemic norms, whether they are of a more internalist or externalist variety. Thus, it is implausible that such general norms imply more specific consciousness-requiring norms. So if there *are* consciousness-requiring norms, it is more plausible to say that they are among the most basic norms than to suppose that they are grounded in other norms.

If the consciousness lover concedes that there is no deeper normative ground for consciousness-requiring norms, and also that there is no plausible external ground for them, then it looks like the only option left is a no-grounding view: consciousness-requiring norms cannot be further explained. If the consciousness lover adopts this position, can anything further be said to arbitrate in a debate between them and a zombie lover? On this option, their disagreement is analogous to a disagreement about what the basic moral norms are. Such basic normative disagreements are notoriously hard to resolve; for example, it seems plausible that they can even survive ideal rational reflection by the disputing parties. Might the disagreement about the rational role of consciousness be like this?

In fact it is not obvious that *all* disagreements about basic epistemic norms are in principle unresolvable in this way. For example, some basic epistemic norms might be knowable a priori. Or more weakly, it might be that there are constraints on what the correct epistemic norms are that are knowable a priori, which at least rule out some proposed norms. Maybe the zombie lover's normative view can be seen to be wrong after sufficient a priori reflection.

Unfortunately, there do not appear to be plausible a priori constraints to settle the dispute. What a priori constraints on epistemic norms are there? In my view, the only plausible candidates have to do with the relationship between the laws of logic (or more broadly logic and probability theory) and the rules of correct reasoning. As Harman (1986) famously emphasized, there is no really straightforward relationship between logical truths and rational norms. Nonetheless, it is plausible to think that rationality is constrained by logic. If the truths of logic are knowable a priori (and it is far from obvious that they are), then perhaps we can know a priori some constraints on the laws of rational reasoning. Even if this is true, however, these logically constrained rational norms are not the ones that are relevant to our debate. The laws of logic constrain the rules governing the rational moves *between* beliefs—what we might call “process norms.” But a zombie can perfectly well reason in a rational way in this sense. We might put this by saying that the zombie is “process rational.” Our dispute is over what might be called “input norms”—norms governing transitions from outside the space of beliefs (e.g., from conscious experiences to beliefs). And these norms are apparently much more remotely related to the kinds of epistemic constraints that can perhaps be known a priori, such as those given by logical rules, or mathematical facts about probability. Or at least, whether inputs into the system have to be conscious or

merely pseudo-conscious appears not to be something to be settled based on such a priori considerations. If this is right, then it is implausible that there is any a priori resolution available in the dispute between consciousness lovers and zombie lovers.

Does this mean that the dispute between a no-grounding consciousness lover and the zombie lover is unresolvable? There are still some considerations remaining that we haven't looked at. In particular, even if consciousness-requiring norms can't be further explained, a case might be made that they are (or are not) *in fact* the correct norms.

On the side of positive arguments for such norms, we might consider cases designed to elicit the intuition that rationality really does require consciousness. First, there is the simple fact that *for us* it seems to be true that if we can't consciously perceive that some state of affairs obtains, then we should refrain from making a non-inferential perceptual judgment about it. For example, if I am trying to find out through perception whether there is a chair in front of me, then, intuitively, unless I have a conscious experience of a chair, any judgment I make is no more justified than a random guess. This intuition seems to hold up, even under the assumption that such "blind guesses" would in fact be reliably correct. This is the moral of blindsight cases: intuitively, if a blindsighted subject has no conscious experience corresponding to a region of space, but can nonetheless make reliable guesses about what is in that region, these guesses are not justified beliefs, unless the subject has some independent reason for thinking that the guesses are reliable. Further, as Smithies emphasizes, these intuitions hold up even if we suppose that we are dealing with a "superblindsight" subject, for whom these judgments come naturally and do not seem like guesswork. According to Smithies, the best explanation of why this is the correct verdict about these cases is that consciousness is strictly necessary for rational judgment.

I personally do not find this conclusion as plausible when I consider a hypothetical functional zombie rather than an individual with localized superblindsight. It is obviously possible to make sense of a functional zombie as a responsible epistemic agent, including in assessing the transitions they make from pseudo-conscious states to beliefs. This suggests that there might be an alternative explanation for why blindsighted guesswork isn't rational. One possibility is that given that we are conscious, consciousness plays a special epistemic role *for us*, so that only conscious states can ground rational belief *for us*. Nonetheless, it could be that for other creatures another kind of internal state—such as a variety of pseudo-consciousness—could play exactly the same role. We can even imagine an alien version of blindsight, in which a region of a pseudo-conscious alien's visual field, although eliciting reliable guesses, is not pseudo-conscious in the right way to allow them to be justified. An alien theorist might explain this by claiming that for aliens, only pseudo-conscious states of a particular kind can serve as reasons for belief. Thus, there might be

an exact symmetry between the norms that the aliens follow and those that we follow, norms that in each case explain why blindsighted guesses, or their alien analogues, are not rational. Consciousness being locally, but not globally, significant in this way, is perfectly consistent with our intuitions about blindsight cases.

Can the zombie lover say anything to motivate this kind of picture over the consciousness lover's vision of consciousness having global epistemic significance? It is at this point that I think reductive materialism, which has not played a role so far in the discussion in this section, might give us a little traction. Suppose we accept the view that if reductionism is true, then consciousness does not mark a deep natural divide between us and the pseudo-conscious aliens. The difference between them and us is merely in how our high-level cognitive architecture is functionally and physically realized, a difference that in no way marks us out as special from the perspective of objective description of the natural facts; a Martian scientist would not find us remarkable relative to our alien counterparts. If there is nonetheless a massive normative difference between us, this is surprising for two reasons.

First, as emphasized by the grounding argument, we intuitively expect a massive normative difference between cases to be reflected in a significant natural difference. That's why if there is no big natural divide between cases of pain and cases of pseudo-pain, this tends to undermine our confidence that pseudo-pain isn't as bad as real pain. Think again about the spectrum of cases linking us with a pseudo-conscious zombie that exists if materialism is true—the trivial natural difference between one case and the next seems insufficient to ground a big normative difference, whether of a practical or epistemic kind. The grounding intuition remains powerful; given that a strong case has *not* been made *in favor* of strong epistemic naturalism, it could be argued that if reductive materialism is true, this puts a burden of proof on the consciousness-loving side.

This is supported by a second observation: if the difference between us and the pseudo-conscious aliens is merely in how our mental states are realized, it is hard to see how this is the right *kind* of difference to ground a deep epistemic divide. Intuitively, rationality or epistemic status are objective notions whose application to a mental state shouldn't depend on such idiosyncratic details, but rather on more relevant features of the context, such as whether the mental state's content stands in the right logical relations to other represented contents. For example, it would be absurd to say that the rationality of a subject's beliefs depends on the color of her hair, because this is an arbitrary characteristic of the subject that has no intelligible relationship to an objective ideal of rationality. Perhaps such a principle of non-arbitrariness could rule against consciousness-requiring norms, if the difference between consciousness and pseudo-consciousness is merely an idiosyncratic difference in the realization of high-level functional structure.

Bearing this in mind, we might describe the state of play as follows: there appear to be no uncontroversial deeper norms that can be appealed to in order

to explain why consciousness-requiring norms might obtain, and attempts to externally ground these norms in the nature of observational concepts or in the intrinsic significance of consciousness do not seem to work. This suggests that such norms will have to be considered brute and inexplicable, but such inexplicable norms seem implausible if consciousness has no deep natural significance: then the difference between conscious and pseudo-conscious beings is merely an arbitrary difference in the way that their high-level architecture is physically realized. The dependence of rationality on such facts offends against the grounding intuition and is puzzling because these realizer facts have nothing to do with the kinds of facts that are paradigmatically associated with rationality, such as facts about the relations between the representational contents of different mental states. I suggest that this puts the burden of proof on those who insist that consciousness *is* necessary for rationality.

Conclusion

Conscious states appear to have very special epistemic properties. According to the grounding argument, they are not uniquely special in this way: if materialism is true, there is no deep natural divide between conscious beings and non-conscious beings, and this means that there is no deep epistemic divide either: an unconscious being could have quasi-conscious states that play much the same epistemic role as ours. We considered a few different ways that a defender of the strong epistemic significance of consciousness could respond to this, focusing on responses that accept reductive materialism, and as a result accept that consciousness lacks strong natural significance. We found these views lacking. It appears that if reductive materialism is true, conscious beings do not occupy as special place in the realm of possibilities as we might have a thought; consciousness is but one among many interesting and unusual ways of configuring physical material.¹³

Notes

1. Although the meaning and theoretical utility of “natural” is typically given by a number of different theoretical roles, not just its role in grounding objective resemblance (see Sider, 2011).
2. See Block (1978) for the related distinction between folk functionalism and empirical functionalism.
3. A qualification: if a Russellian Monist version of reductionism is true (see Chalmers, 2003; Stoljar, 2006), then there may be a sense in which knowing the real definition of consciousness in terms of physical/functional concepts won't amount to full knowledge of which property it is. A full argument for deflationism would need to consider this case.
4. Notoriously, this is a view rejected by certain theorists, like Davidson (1986), who thought that “only a belief can justify a belief.” We will not consider this view here.
5. This latter view could be regarded as a kind of normative version of a dispositional higher-order thought view (e.g., Carruthers, 2003): what makes a state conscious is that it is immediately available for justified self-ascription.
6. On a weaker interpretation, the role of norms is not in individuating beliefs and desires, but only in making sense of the epistemology of belief ascription: we need to use the principle

- of charity to decide between competing attitude ascriptions that are consistent with the subject's overt behavior (see Kriegel, 2011, [chaps. 1 and 4](#)).
7. See Boyd (1988) for a classic defense of normative reductionism in the case of moral norms.
 8. Note that the point is not that realism about norms is incompatible with reductive materialism, as I am understanding it here. The grounding of consciousness facts in functional/material facts is compatible with primitive normative facts existing—it will just have to be that the normative facts are not part of the ground of the consciousness facts.
 9. Even nonreductivists about consciousness might want to reject nonreductionist realist normativism for similar reasons. For example, most property dualists regard phenomenal properties as natural properties, at least in the sense that they might be subject to natural laws in much the way fundamental physical properties are. Presumably, this is incompatible with them being individuated in irreducibly normative terms. Related to this, the idea that there is a scientific project of figuring out what consciousness is in natural terms arguably depends on the assumption that it is a wholly natural phenomenon. If it isn't, then such a project is, arguably, misconceived in much the way that a scientific investigation into the nature of objective moral goodness would be.
 10. See Parfit (2011), part 6, for a helpful discussion of the difficulties with normative reductionism in the case of practical normativity.
 11. For classic statements of the error theory and expressivism, respectively, see Mackie (1977) and Hare (1952) (both focus on practical, rather than epistemic norms). Field (2009) argues for a version of expressivism about epistemic norms.
 12. Despite this objection, psychological expressivism has been endorsed in various forms by a number of authors, most notably by Wittgenstein in his remarks on pain ascription (Wittgenstein, 1953). A notable recent attempt to develop the view is Hellie (manuscript).
 13. Thanks to Wes Holliday, Uriah Kriegel, and Declan Smithies for helpful feedback and discussion of earlier drafts.

References

- Armstrong, D. (1978). *Universals and Scientific Realism: A Theory of Universals*. Vol. 2. Cambridge: Cambridge University Press.
- Block, N. (1978). Troubles with functionalism. *Minnesota Studies in the Philosophy of Science* 9: 261–325.
- Block, N. (2002). The harder problem of consciousness. *Journal of Philosophy* 99 (8): 391–425.
- Boyd, R. (1988). How to be a moral realist. In G. Sayre-McCord (ed.), *Moral Realism*. Ithaca, NY: Cornell University Press. pp. 181–228
- Campbell, J. (2002). *Reference and Consciousness*. Oxford: Oxford University Press.
- Carruthers, P. (2003). *Phenomenal consciousness: A naturalistic theory*. Cambridge: Cambridge University Press.
- Chalmers, David J. (2003). Consciousness and its place in nature. In Stephen P. Stich & Ted A. Warfield (eds.), *Blackwell Guide to the Philosophy of Mind*, pp.102–143. Malden, MA: Blackwell.
- Dainton, B. (2008). *The Phenomenal Self*. New York: Oxford University Press.
- Davidson, D. (1973). Radical interpretation. *Dialectica* 27 (1): 314–328.
- Davidson, D. (1986). A coherence theory of truth and knowledge. In Ernest LePore (ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*. Oxford: Basil Blackwell.
- Dennett, D. (1971). Intentional systems. *Journal of Philosophy* 68: 87–106.
- Field, H. (2009). Epistemology without metaphysics. *Philosophical Studies* 143 (2): 249–290.
- Goodman, N. (1983). *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.
- Hare, R.M. (1952). *The Language of Morals*. Oxford: Clarendon.
- Harman, G. (1986). *Change in View: Principles of Reasoning*. Cambridge, MA: MIT Press.
- Hellie, B. (unpublished manuscript). “Out of this World.”
- Kriegel, U. (2011). *The Sources of Intentionality*. New York: Oxford University Press.

- Lee, G. (Forthcoming). Alien subjectivity and the importance of consciousness. In A. Pautz & D. Stoljar (eds.), *Themes from Block*. Cambridge, MA: MIT Press.
- Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy* 61: 343–377.
- Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. Harmondsworth; NY: Penguin.
- McDowell, J. (1985). Functionalism and anomalous monism. In Brian P. McLaughlin & Ernest LePore (eds.), *Action and Events*. Oxford: Blackwell.
- McLaughlin, B. P. (2003). A naturalist-phenomenal realist response to Block's harder problem. *Philosophical Issues* 13: 163–204.
- Parfit, D. (2011). *On What Matters*. Vol 2. Oxford: Oxford University Press.
- Peacocke, C. (1992). *A Study of Concepts*. Cambridge, MA: MIT Press.
- Peacocke, C. (2004). *The Realm of Reason*. New York: Oxford University Press.
- Prior, A. (1967). The runabout inference ticket. In Peter Strawson (ed.), *Philosophical Logic*. London: Oxford University Press.
- Searle, J. (1989). Consciousness, unconsciousness, and intentionality. *Philosophical Topics* 17 (1): 193–209.
- Sider, T. (2011). *Writing the Book of the World*. Oxford: New York : Oxford University Press.
- Smithies, D. (2012a). The mental lives of zombies. *Philosophical Perspectives* 26: 343–372.
- Smithies, D. (2012b). A simple theory of introspection. In D. Smithies & D. Stoljar (eds.), *Introspection and Consciousness*, 259–293. New York: Oxford University Press. p. 259.
- Stoljar, D. (2006). *Ignorance and Imagination: The Epistemic Origin of the Problem of Consciousness*. Oxford: Oxford University Press.
- Taylor, B. (1993). On natural properties in metaphysics. *Mind* 102: 81–100.
- Wittgenstein, L. (1953). *Philosophical Investigations*. New York: Macmillan.

Suggestions for Further Reading

Horgan, Terry, and John Tienson. 2002. "The Intentionality of Phenomenology and the Phenomenology of Intentionality." In *Philosophy of Mind: Classical and Contemporary Readings*, ed. D.J. Chalmers, 520-533. Oxford: Oxford University Press.

This paper argues against the separation of the mind into a sensory, phenomenal domain and a cognitive, intentional domain—and in favor of an "inseparatist" picture in which the sensory is inherently intentional and the cognitive is inherently phenomenal.

Ryle, Gilbert. 1949. *The Concept of Mind*. London: Routledge, 2009.

This influential book systemically argues that our very grasp of the mental domain is founded on our grasp of dispositions to behave in certain ways, not on anything like direct acquaintance with inner consciousness.

Siewert, Charles. 1998. *The Significance of Consciousness*. Princeton, NJ: Princeton University Press.

This book presents a thoroughly first-personal approach to consciousness and its centrality to mental life, making influential arguments that such phenomena as intentionality, cognition, and self-knowledge are shot through with phenomenality.

Strawson, Galen. 1994. *Mental Reality*. Cambridge, MA: MIT Press.

This book argues that the mental domain is essentially the domain of conscious experience, and any connection to behavioral dispositions is limited, indirect, and accidental.

Supplemental Guide to Further Controversies

The previous sections in this volume present five of the most central controversies in the philosophy of mind. Listed below are three additional important controversies for the curious reader, with lists to the essential readings for each controversy.

Perception and Perceptual Representation

- Ayer, A. J. 1956. *The Problem of Knowledge*. London: Macmillan.
- Byrne, Alex. 2001. "Intentionalism Defended." *Philosophical Review* 110: 199–240.
- McDowell, John. 1994. "The Content of Perceptual Experience." *Philosophical Quarterly* 44: 190–205.
- Siegel, Susanna. 2010. *The Contents of Visual Experience*. New York: Oxford University Press.

Cognitive Phenomenology

- Carruthers, Peter, and Bénédicte Veillet. 2011. "The Case against Cognitive Phenomenology." In *Cognitive Phenomenology*, ed. T. Bayne and M. Montague, 35–56. Oxford: Oxford University Press.
- Chudnoff, Elijah. 2011. "What Intuitions Are Like." *Philosophy and Phenomenological Research* 82: 625–654.
- Moore, G. E. 1953. "Propositions." In *Some Main Problems of Philosophy*. London: Routledge.
- Pitt, David. 2004. "The Phenomenology of Cognition—or What Is It Like to Believe That P?" *Philosophy and Phenomenological Research* 69: 1–36.

Mental Causation

- Block, Ned. 1990. "Can the Mind Change the World?" In *Meaning and Method: Essays in Honor of Hilary Putnam*, ed. G. Boolos. Cambridge: Cambridge University Press.
- Kim, Jaegwon. 1989. "Mechanism, Purpose, and Explanatory Exclusion." *Philosophical Perspectives* 3: 77–108.
- Stich, Steven. 1978. "Autonomous Psychology and the Belief-Desire Thesis." *Monist* 61: 573–591.
- Yablo, Steven. 1992. "Mental Causation." *Philosophical Review* 101: 245–280.

This page intentionally left blank

Index

- Armstrong, D. M. 23, 37–8, 55, 65–7, 224, 244
- Block, N. 89, 133–4, 181, 183, 194, 218, 220, 227, 243–5, 247
- Brentano, F. C. 117, 170, 177–8
- Byrne, A. 67, 131, 133, 153, 155–6, 247
- Chalmers, D. J. 29–31, 33, 35–9, 42–8, 55–7, 59–64, 66–8, 70, 106–7, 109, 115, 131, 133, 136, 143–4, 157, 176, 178–9, 181, 194, 219–30, 243–4, 246
- consciousness 1–2, 5, 17–18, 22, 26, 38, 40–2, 47–9, 54–5, 63, 113, 154, 168, 199–200, 203, 206–8, 211, 213–16, 218–20, 229, 234, 240–4; importance/significance of 10–12, 201, 217, 222–5, 227–31, 235–9, 243; intentionalism about 6–7; nature of 1, 19, 26, 58–9, 230, 232; phenomenal 10–11, 43, 49, 60, 62, 114, 136, 181, 200–1, 204, 216, 226, 233; problem of 114–15, 129, 131; pseudo- 11, 226–7, 229–32, 234, 236, 241–2; truths about 33–5, 58–9; value of 10–11, 220
- Crane, T. 115, 132–3, 136, 157–8
- Dennett, D. 175, 178, 214, 219–20, 233, 244
- Descartes, R. 43, 50, 63–4, 67–9, 92
- Dretske, F. 13, 115, 133, 143, 153, 157, 166, 174, 178, 180, 191, 194, 196, 218, 220
- dualism 2–4, 11–12, 17–9, 27, 29, 34, 38, 178, 194, 219, 229, 244
- embodied cognition 1, 73–5, 77–8, 84, 86–7, 98, 100; body's role in 83; characterization/definition of 5–6, 73, 79–83, 85, 91, 99, 102; research in 5, 76, 81, 88–9, 92–4, 103–6
- Fodor, J. 93, 107, 143, 157, 165, 173, 178, 180–6, 191, 193–6
- Goldie, P. 116, 123, 132–3, 138, 145, 147, 154–5, 157
- Goldman, A. I. 5–6, 13, 73–4, 79–92, 102–3, 105, 107, 109
- Harman, G. 129, 133, 156–8, 169, 178, 240, 244
- Intentionalism 7–8, 114, 135–42, 144–8, 150–6

- intentionality 1, 8, 12, 114–15, 122, 131, 170–1, 173, 177–8, 183, 190, 193, 236; naturalistic theories of 6–7, 9, 131, 165–7; of emotions 121, 132; of moods 7–8, 113, 115–26, 130, 132–3, 147–53; phenomenal 154, 174, 206–7, 219
- intrinsic property 4–5, 23, 41, 48–9, 64, 70; absolutely 35, 38, 40, 49–63, 65–7
- Jackson, F. C. 37, 38, 63–4, 66, 68–9, 153, 157, 195
- Kant, I. 37, 39–40, 50, 53–4, 64–6, 68–9, 177
- Kind, A. 7, 131–4, 149–51, 156–7
- Kriegel, U. 36, 67, 155–7, 173–4, 176, 178–9, 188, 191–5, 219–20, 244
- Lee, G. 11, 217, 225, 229, 245
- Leibniz, G. W. 3–4, 13, 49–56, 60, 64–5, 68
- Lewis, D. K. 23–5, 29–30, 33, 37–9, 54, 66, 68–9, 220, 224, 245
- Locke, J. 19–23, 37, 39, 52, 54–5, 65–6, 68–9
- Lycan, W. G. 125, 134, 136, 157, 218, 220
- McGinn, C. 56, 62, 66–7, 69, 132, 134, 167, 178–9, 196
- materialism 1–3, 6–7, 11–13, 17–28, 30–4, 36, 38, 41–2, 45, 48; nonreductive 12, 229; reductive 12, 217, 222–4, 226, 228–30, 234, 236–7, 242–3
- Mendelovici, A. 7–8, 120–1, 127–8, 131–3, 155, 157
- mood: and emotion 114, 116–17, 131–2, 135–8, 146–8, 152–3, 155–6; intentionality of 7–8, 113, 115–26, 130, 132–3, 147–53; phenomenal character of 6–7, 126–9, 132, 136, 146–7, 153
- naturalism 6, 13, 165, 172, 181, 193, 230, 235–6, 242
- Pereboom, D. 4–5, 35–9, 63–4, 66, 69
- phenomenology 6–7, 56, 115, 119, 125–8, 130, 139–40, 145, 152, 206, 221, 246–7
- physicalism 36, 41, 43, 48–9, 55–6, 58–60, 62–4, 66–7, 114, 178, 218
- Prinz, J. 94, 108, 131–2, 134, 139, 157
- representation 7–8, 52, 59–62, 74–5, 77–9, 83, 88, 92, 95, 98, 119, 135–42, 145, 151, 154–5, 211, 219, 247; format of 5–6, 79–80, 82, 89, 102, 104–5; mental 1, 12, 84, 104, 136, 143, 161–3, 165–8, 171–7, 180–8, 191–3, 196, 201, 223; objective 9, 161–8, 172–5, 177–8, 180, 191–2; of the body 80, 85, 89, 93–4, 97, 100–2, 105–7, 122–3, 139, 146–7; subjective 9, 161–75, 177–8, 187–8, 190–3; unbound 8, 121–2, 125, 128, 132–3, 148–50
- representationalism 7, 75, 93, 99, 113–16, 118–19, 121, 124, 126–7, 129–32, 207, 219
- Rupert, R. 9, 89–90, 169, 184, 188, 190–1, 193, 195
- Russell, B. 13, 27–9, 31, 37–9, 41, 63, 69–70
- Russellian monism 2, 4–5, 17, 19, 23, 27, 31, 33–4, 36–8, 40–3, 45–6, 48–9, 54–6, 58–60, 62–3, 243
- Seager, W. 7, 13, 138, 147, 157
- Searle, J. R. 119, 134, 173–4, 178–9, 218–19, 221, 236, 245
- Shapiro, L. 6, 75–6, 78–9, 89–91, 103–7, 178
- Shoemaker, S. 54, 66–7, 69, 144, 157
- Siewert, C. 10, 131, 134, 178–9, 218–21, 223, 236, 246
- Stoljar, D. 4, 34, 36, 38–42, 48–9, 54, 56–9, 63, 66–7, 69–70, 155–6, 220–1, 243, 245
- Strawson, G. 55–6, 60, 66–7, 69, 176–9
- Tye, M. 123, 132–4, 136, 138, 141–2, 145–7, 153–7, 218, 221
- unconscious 10, 27, 200, 217, 223–5, 227–31, 236, 243