

# Vagueness and zombies: why ‘phenomenally conscious’ has no borderline cases

Jonathan A. Simon<sup>1</sup>

Published online: 24 September 2016  
© Springer Science+Business Media Dordrecht 2016

**Abstract** I argue that there can be no such thing as a borderline case of the predicate ‘phenomenally conscious’: for any given creature at any given time, it cannot be vague whether that creature is phenomenally conscious at that time. I first defend the Positive Characterization Thesis, which says that for any borderline case of any predicate there is a positive characterization of that case that can show any sufficiently competent speaker what makes it a borderline case. I then appeal to the familiar claim that zombies are conceivable, and I argue that this claim entails that there can be no positive characterizations of borderline cases of ‘phenomenally conscious’. By the Positive Characterization Thesis, it follows that ‘phenomenally conscious’ can not have any borderline cases.

**Keywords** Phenomenal consciousness · Vagueness · Explanatory gap · Conceivability argument · Knowledge argument

## 1 Introduction

I claim that ‘phenomenally conscious’—the predicate that applies to any entity that there is *something it is like* to be<sup>1</sup>—is not vague. Intuitively, asking whether someone is phenomenally conscious is like asking whether the light is on. The light

---

<sup>1</sup> For discussion of this locution see Nagel (1974), Block (2002). The predicate at issue should be read as implicitly time-stamped. I do not deny that it can be indeterminate whether someone counts as having been conscious for *most* of some interval. Also, I do not deny that other phenomenal terms, like ‘pain’ can

---

✉ Jonathan A. Simon  
jas741@nyu.edu

<sup>1</sup> NYU Department of Philosophy, Center for Mind, Brain and Consciousness, New York NY, USA

may be brighter or dimmer, but if there is any luminescence at all, then it is on.<sup>2</sup> Experience may be more or less intense, more or less attentive, closer to waking or closer to dreaming, but if at some specific moment it is not completely, totally, absolutely dark inside, then at that moment there is some phenomenal consciousness.

In support of this first claim I make a second, more general one: that when something is a borderline case of a predicate, there is a *positive characterization* of that borderline case—a way of filling in the details of the case that will show any sufficiently competent speaker what makes it a borderline case, and accordingly put any such speaker in a position to see why it is a borderline case. I'll call the thesis codifying this idea the Positive Characterization Thesis.

In this paper, I will defend the Positive Characterization Thesis (showing it to be a development of the idea that vagueness is a matter of how we ought to use our terms and concepts), and then I will argue for the conditional claim that if the Positive Characterization Thesis is true, and if also zombies are conceivable, then 'phenomenally conscious' is not vague.

There are several reasons to care about whether I am correct. First, the Positive Characterization Thesis has implications concerning what vagueness consists in, and whether there can be an effective test of it, matters that have been contested in the recent literature.<sup>3</sup>

Second, there is the interest in itself of the pull we feel to think of phenomenal consciousness as a kind of 'inner light' that must be either on or off. Many have admitted to feeling the force of this intuition,<sup>4</sup> but the only sustained development of it in the literature, due to Michael Antony,<sup>5</sup> though full of valuable insights, is vulnerable to serious objections.<sup>6</sup> Indeed, there is a basic tension in squaring this

---

Footnote 1 continued

be vague: it can be vague whether some experience is a pain or a tickle. But this presupposes an experience of some sort or other. For this reason, the so-called *phenomenal sorites* is no evidence of vagueness in 'phenomenally conscious'. See Sebastian (2011) for discussion.

<sup>2</sup> Others who make this analogy include McGinn (1996), Papineau (1993). As it happens the analogy is imperfect, because some will allege that indeterminacy arises at the quantum level (though this is a subtle matter). Note that with light it is natural to speak in terms of degree: we might say that a light is on, provided it is luminescent to some degree. I am prepared to allow that consciousness, like light, comes in degrees of intensity. What I deny is that attributions of consciousness therefore come in degrees of truth. Accordingly, I allow that 'phenomenally conscious' is a minimum standard absolute adjective, but I deny that it is a relative gradable adjective (Kennedy 2007). Another predicate behaving this way is 'voluminous'.

<sup>3</sup> For example Antony (2006a, b), Eklund (2005), Greenough (2003), Schiffer (2003), Smith (2005), Weatherston (2010), and Wright (2001).

<sup>4</sup> See for example Chalmers (2013), McGinn (1996), Searle (1992), Sebastian (2011), Papineau (1993), Tye (1996), Unger (2004). Those who acknowledge the intuition but have argued that it is misleading include Brogaard (2010), Dennett (1998), Deutsch (2005), Papineau (2002), Tye (1996), Unger (1988).

<sup>5</sup> Antony (2006a, b, 2008). But see also Goff (2013) and Sebastian (2011).

<sup>6</sup> My argument owes much to Antony's. Like Antony, I argue that there is a necessary condition on vagueness which 'phenomenally consciousness' does not meet because of its peculiar semantics. But the devil is in the details. The necessary condition that Antony formulates, as he actually states it (his C1-C4), is too weak, requiring only that a vague concept includes some constraints common to instances,

intuition with the fact that ‘phenomenally conscious’ occurs in folk psychology and applies to messy macroscopic entities and is difficult to convey to undergraduates. How could such a predicate fail to be vague?<sup>7</sup>

Third, at stake is what we should think about the status of creatures like fish, of developing fetuses, and of persons in vegetative or minimally conscious states. Many suppose the moral status of these beings to be a function of whether they are phenomenally conscious. Is it an option to think of them as having some kind of borderline moral status? This probably depends on whether they are borderline cases of ‘phenomenally conscious’.

Fourth, the question of vagueness is relevant to the metaphysics of consciousness. The leading analyses of phenomenal consciousness are couched in vague terms, so, as Antony points out,<sup>8</sup> if ‘phenomenally conscious’ is not vague, no leading analysis can be more than approximately correct. But more is true. Though it will be beyond my scope to consider this matter here, materialists *rely* in various ways on the vagueness of ‘phenomenally conscious’. To put the matter intuitively: the natural world is a world of gradual variation: how could ‘phenomenally conscious’ fail to be vague, if phenomenal consciousness is a fully integrated part of this gradual natural world?<sup>9</sup>

The challenge to materialism that arises here is distinctive because, though it depends on the Positive Characterization Thesis, it is independent of the more problematic steps in arguments such as Jackson’s Knowledge argument or

---

Footnote 6 continued

borderline cases, and non-instances—satisfied by the constraint that, e.g., the things in question *exist*. As he seems to intend it however the constraint is given by a dimension of variation, such that variation in this dimension explains whether something is an instance, a borderline case, or a non-instance. But this is quite strong. Brogaard (2010) argues that it is too strong, and would rule out that ‘bald’ is vague, and it would likely also rule out that Weatherson (2010)’s ‘tall-179’ is vague. In any event it calls for substantive defense. Another problem for Antony is that his specification of what is peculiar about the semantics of ‘phenomenally conscious’ involves a speculative empirical hypothesis about the mechanics of thought: that the complex mental file associated with the concept ‘phenomenally conscious’ does not contain any material ‘elements’. For Antony’s argument to succeed, this must amount to more than simply the claim that there are no a priori materialist analyses of consciousness. For example, you might think it to be conceptually necessary that conscious beings are functionally complex, even if you think that zombies are conceivable. But then a spectrum of functional complexity could allow ‘phenomenally conscious’ to satisfy Antony’s condition. See chapter one of my 2012 for a sustained critique of Antony’s argument.

<sup>7</sup> Brogaard (2010), Tye (1996, 2000). Weatherson (2003) argues that ‘phenomenally conscious’ must be vague, since ‘David Chalmers’ is a vague name (think: problem of the many), but ‘David Chalmers is phenomenally conscious’ is determinately true. But there is no tension here: it suffices that all of the precisifications of ‘David Chalmers’ are phenomenally conscious. This needn’t entail that there are multiple conscious beings, any more than it entails that there are multiple David Chalmerses. For example it might be that there is some *state* common to each precisification which suffices for phenomenal consciousness. But see Unger (2004).

<sup>8</sup> Antony (2006a, b, 2008).

<sup>9</sup> Consider for example the metasemantic challenge: what could privilege a single candidate referent as more eligible than all of the others? My argument shows that ‘phenomenally conscious’ is non-vague but without providing materialists with a recipe for selecting a most eligible candidate. As Papineau (2002) points out, materialists already encounter a metasemantic problem—one that Balog (ms) calls the *Hardest Problem of Consciousness*—even if we allow that ‘phenomenally conscious’ is vague. But if my argument here succeeds, this problem gets even harder.

Chalmers' Conceivability argument, and accordingly it is immune to many of the leading objections to those arguments.<sup>10</sup>

Thus, much hinges on whether 'phenomenally conscious' is vague. Here is how the paper will proceed. In Sect. 2, I will introduce and defend the Positive Characterization Thesis. In Sects. 3–5 I defend the conditional claim that *if* the Positive Characterization Thesis is correct then *if* zombies are conceivable, 'phenomenally conscious' is not vague. In Sect. 3 I make precise the claim that zombies are conceivable, and I identify a few further claims that those who endorse it should also accept. In Sect. 4 I argue, drawing on Sect. 3, that descriptions that conceivably describe determinate zombies cannot be positive characterizations of borderline cases of 'phenomenally conscious'. Then in Sect. 5 I argue by process of elimination that no description that doesn't conceivably describe a determinate zombie can be a positive characterization of a borderline case of 'phenomenally conscious' either. Putting the immediate conclusions of Sects. 2–5 together, I conclude that 'phenomenally conscious' is not vague.<sup>11</sup>

## 2 The Positive Characterization Thesis

There is more to vagueness than indeterminacy. Suppose Kripke (1982) is right that most terms, including 'plus', are indeterminate in their application. Most of us would not take it to follow that 'plus' is vague. Some take the lesson to be that vagueness must involve *soriticality*: a vague predicate must be a *tolerant* predicate in the sense of Wright (1975), one that gives rise to soritical reasoning. However, others point to cases that appear to be vague without being soritical, for example Weatherston (2010)'s 'few children for an academic'.<sup>12</sup>

But whether soriticality is coextensive with vagueness or not, there is more to say about the intuition that 'plus' wouldn't be vague even if it were indeterminate. The intuition ultimately turns, I contend, on the idea that vagueness is something to which ordinary competent speakers can in principle be sensitive. There is no way we could reasonably expect ordinary competent speakers to be sensitive to the indeterminacy of 'plus'.

<sup>10</sup> Broadly speaking, I have in mind the objections to Chalmers and Jackson that target those thinkers' modal rationalism: the view that there is a direct constitutive connection between the facts about what is a priori or conceivable on the one hand and what is necessary or possible on the other. The only constitutive connection in the challenge I envision is a connection between the facts about semantic competence on the one hand and facts about vagueness on the other. The rest of the work is done by independently motivated metaphysical premises. See my 2012 for further discussion.

<sup>11</sup> Where possible, I will formulate the claims of this paper in terms of terms, predicates and descriptions rather than concepts and propositions. This allows me to retain neutrality about what concepts and propositions are (for example, whether they are vague, as in Lewis (1975)), and it makes certain formulations (e.g. about the role of context) more natural. But I take it that much of what I say will carry over. For example, I intend for my argument here to translate into an argument that the phenomenal concept PHENOMENALLY CONSCIOUS is not vague. Compare my disclaimer here to similar remarks in Boghossian (2003a), deRosset (2013).

<sup>12</sup> See Gaifman (2010) for further examples. See Alston (1964) for the pluralist compromise: he distinguishes between *degree* and *combinatorial* vagueness.

Here, I'll propose an abstract constraint on what this sensitivity involves. This will fall short of a fully specific theory of what vagueness consists in, but it will have some significant consequences nevertheless.

I'll focus on vagueness in predicates. I mean to use 'borderline case' as a term of art, not equivalent to 'indeterminate case'. I take it that  $x$  is a borderline case of 'P' just in case it is vague whether  $P(x)$ . If Kripke is correct, then the predicate '67 plus  $58 = x$ ' may be indeterminate, but it is not vague.

For reasons I will make clear below, I take the intuition about sensitivity and competence on which my proposal turns to be best expressed as a claim about which patterns of usage of the terms in question are *appropriate* in which circumstances, or as I shall say, a claim about the *norms of usage*. But as I will explain, there is no deep tension between my approach and one that speaks of which dispositions are concomitant of full competence, i.e. mastery. The normative element on which I mean to focus is embedded in the notion of full competence.

I will suggest that a certain class of norms is distinctive of borderline cases: that is, the norms in this class are jointly in force for all (and perhaps only) borderline cases, and they specify which semantic, psychological or epistemic attitudes it is appropriate or inappropriate for us to adopt when confronted with such cases.<sup>13</sup>

There are deep philosophical questions about the logic of vagueness. These questions are reflected in questions about the usage norms concerning vagueness, broadly construed. Ought we withhold judgment while affirming classical tautologies as well as classical inference rules? Then epistemicism may be correct. Ought we withhold judgment while affirming classical tautologies but rejecting some classical inference rules? Then supervaluationism gives us the right picture. Ought we reject even some classical theorems? Then some many-valued approach may be best.

I don't claim that the normative facts are what explain the facts about logic.<sup>14</sup> It may be that the facts about logic (inter alia) explain the normative facts. What I am going to suggest is a specific manner in which these normative facts are a part of what vagueness consists in: that what it is for a given application of a term to be a borderline case is (at least in part) for a certain group of norms to correctly characterize it.<sup>15</sup>

Why endorse a claim like that? For one thing, it suggests an explanation of the intuition that vagueness is something to which sufficiently competent speakers can in principle be sensitive. For another thing, if we acknowledge that distinctive usage norms concerning vagueness exist anyway, then it is a matter of theoretical economy to suppose that these norms play some role in explaining vagueness; that is, in specifying what it consists in.

<sup>13</sup> I remain neutral about whether there are sui generis semantic norms (in the sense of Boghossian (2003b), Glüer and Wikfoss 2015), or whether all of the norms at issue ultimately derive from epistemic, psychological, pragmatic or practical norms.

<sup>14</sup> See for example Boghossian (2003a), Brandom (1994) for different versions of this idea.

<sup>15</sup> Thus my claim is that if the logic comes first the logic gives rise to vagueness in virtue of giving rise to the norms.

Some recent authors (most notably Eklund (2005), but also arguably Antony (2006), Dummett (1975), Fine (1975), Greenough (2003), Smith (2005), Schiffer (2003), and Wright (2001)) can be construed as agreeing with something in the vicinity of this general claim, though many opt for formulations in terms of the dispositions of competent speakers, rather than usage norms—an issue I will return to shortly. Whichever formulation we opt for, there is a question of precisely how much of an explanation we may hope for. May we hope, as I suggest, for explanation of what makes a given borderline case borderline, or may we only hope for explanation of what makes a given predicate vague, generally speaking? Some, like Greenough (2003) and Eklund (2005), suppose only that there are constraints on competent reasoning with a vague predicate in general. Eklund, for example, argues that we must accept the principle:

(Eklund's Principle) Fully competent speakers are disposed to accept that whereas large enough differences in F's parameter of application sometimes matter to the justice with which it is applied, some small enough difference never thus matters.

Weatherson (2010) (to whom this formulation of Eklund's Principle is due) points out that the principle contains a two-way scope ambiguity. Is the claim **Strong**: that there is an actual parameter of application such that all competent speakers are disposed to accept the relevant claim about *it*, or is the claim **Weak**: that speakers are disposed to accept the existential claim that there is some parameter of application or other such that the relevant claim holds true about it?

One critical difficulty for **Strong** to which I will return below is the problem of error and ignorance: competent speakers may disagree over which are the true parameters of application (as they also might over which cases are borderline). A difficulty for **Weak** is that it requires that competent speakers have thoughts about parameters of application as such, which over-intellectualizes.<sup>16</sup>

My primary concern, however, is that neither proposal offers an analysis of borderlineness. Perhaps there are borderline cases in the middle of any properly soritical series, but more must be said for this insight to lead to a genuine analysis.<sup>17</sup>

But why would there be a use-normative account of what sorites series consist in, if there were not also such an account of what borderline cases consist in; of what makes a given case a borderline case? Must we appeal to entirely separate considerations to give an account of borderlineness?

<sup>16</sup> What about **Medium**: for each competent speaker there is some parameter such that that speaker is disposed to accept the relevant claim about that speaker's favored parameter? This may avoid the problem of disagreement, but at the cost of yielding false positives.

<sup>17</sup> Antony's proposal comes closer than Eklund's to offering such an account: he speaks of borderline cases as individuals with respect to which a certain class of dispositions distinctive of borderline cases, *v-dispositions*, is manifested, when those cases are thought about under suitable 'individual conceptions'. This is very much on the right track, but Antony falls short in spelling out what makes for a suitable individual conception, as I argue at length elsewhere (and sketch in note 6 above).

I claim that a use-normative account of borderlineness is indeed available. To offer such an account, we must identify norms of usage specifying what attitudes it is appropriate (or inappropriate) to take when confronted with a borderline case.

I take it that usage norms for borderline cases may encompass psychological norms (e.g., that one ought not use classical probabilistic principles when reasoning about the cases in question<sup>18</sup>) and epistemic norms (e.g., that one ought to be particularly cautious about the inferences one draws involving the cases in question<sup>19</sup>), as well as pragmatic and contextual norms (e.g., that one had ought to pay attention to shifts in the context when reasoning about the cases in question<sup>20</sup>). And since we do not have unmediated vagueness detectors, these norms should be couched in terms of beliefs concerning the cases in question.

Thus what we seek are usage norms of the form:

It is appropriate for those who believe that  $R(x)$  in a given context to  $\Phi(x, 'P')$  in that context.<sup>21</sup>

Where 'R' is a predicate and to  $\Phi(x, 'P')$  is to take the vagueness-appropriate attitudinal stance towards the application of 'P' to  $x$ . ' $\Phi$ ' is a placeholder. My aim in the below will not be to arbitrate between rival accounts of what attitudes are appropriate in the presence of borderline cases. I suppose only that there exist a set of such attitudes. I suspect that in addition these norms are proprietary in the sense that, *if it is appropriate for those who believe that  $R(x)$  to  $\Phi(x, 'P')$ , then  $R$ s are borderline cases of 'P'*, but I will not rely on this assumption below (though I will consider it further).<sup>22</sup>

But the existence of such norms does not yet give us a use-normative *account* of vagueness, not even if the norms in question are proprietary. For that, we need some of the norms of the above form to help *constitutively explain* the borderlineness of the case in question. That is, we need  $R$ s such that it is (in part<sup>23</sup>) *because* it is appropriate to take the  $\Phi$ -attitudes towards believed  $R$ s (in a given context), that satisfying 'R' secures that something is a borderline case (in that context). I will call usage norms that play this explanatory role *explanatory usage norms*.

I claim that such explanations are available for every borderline case of every predicate. If there is a reason that you think a case is borderline, there will be an explanatory usage norm accounting for the feature of the case you are responding to. For example, it is because it is appropriate to adopt the relevant attitudes concerning

<sup>18</sup> Schiffer (2003).

<sup>19</sup> See Williamson (1994).

<sup>20</sup> Fara (2000), Gaifman (2010), Raffman (1996), Shapiro (2008).

<sup>21</sup> We might also look for norms of inappropriateness, obligation, impermissibility, and so on, that are characteristic of borderline cases. For reasons of economy I will take norms of appropriateness to be representative.

<sup>22</sup> Should we only be talking about the case of those who are suitably *guided* by their belief that  $R(x)$  in adopting the  $\Phi(x, 'P')$  attitudes? I mean to remain neutral. Do the standards at issue here admit of accidental compliance? I invite those who think not to think of a guidance (or non-deviant causation) condition as built into the specification of the norms I consider. Nothing I say below precludes one, but for those who are unmoved, nothing I say below requires one.

<sup>23</sup> If the norms are proprietary then the explanation may be complete. Otherwise it will only be partial.

the predicate ‘tall’ toward things one believes to be 5’10” (in some suitable contexts), that satisfying ‘5’10” secures that something is a borderline case of ‘tall’ (in those contexts). It is because it is appropriate to adopt the relevant attitudes concerning ‘rich’ toward people one believes to earn roughly \$200,000 a year (in some contexts) that satisfying ‘earning \$200,000 a year’ secures that something is a borderline case of ‘rich’ (in those contexts). It is because it is appropriate to adopt the relevant attitudes concerning the predicate ‘orange’ toward things one believes to be  $\text{rgb}(255,69,0)$  (a shade on the borderline between orange and red), or to be *this shade* (a perceptual demonstrative for  $\text{rgb}(255,69,0)$ ) that satisfying those predicates secures that something is a borderline case of ‘orange’.<sup>24</sup> I will have more to say about difficult cases below.

Note that it is not an explanatory usage norm if ‘R’ says that its instances are borderline cases of the predicate in question. Suppose ‘R’ = ‘borderline case of ‘P’’. If the relevant norm were explanatory, this would mean that it is because it is appropriate to adopt the  $\Phi$ -attitudes towards things one believes to be borderline cases of ‘P’, that satisfying ‘borderline case of ‘P’” secures that something is a borderline case of ‘P’. But obviously there is no such explanatory connection: disquotational principles exhaustively explain why satisfying ‘borderline case of ‘P’ secures that something is a borderline case of ‘P’. Such an ‘R’ may figure in a usage norm, but this norm will not be an explanatory one. Say that an ‘R’ that figures in an explanatory usage norm (in my strict sense) is a *positive characterization of a borderline case*. Such an ‘R’ will not merely say that vagueness is afoot; it will show it, as ‘5’10” does in the case of ‘tall’. The Positive Characterization Thesis says that every borderline case of every predicate has a positive characterization.

This is the Positive Characterization Thesis, officially stated:

(Positive Characterization Thesis) For any entity  $x$ , context  $c$ , and predicate ‘P’, if  $x$  is a borderline case of ‘P’ in  $c$ , there is a description, ‘R’, such that ‘R( $x$ )’ is true in  $c$ , and such that it is an explanatory usage norm that in  $c$ , it is appropriate for those who believe that R( $x$ ) to  $\Phi(x, \text{‘P’})$ .

The Positive Characterization Thesis thus is an explanatory thesis: it states that for every borderline case there is an explanation of why that case is borderline in terms of norms that apply to that case in light of what else is true about it (i.e., in light of its satisfying suitable ‘R’). Without ruling on many of the subtler questions about the correct logic, semantics, epistemology or psychology of vagueness, this thesis outlines a general pattern of explanation of vagueness which accounts for our intuitions concerning why vagueness is something to which we can be sensitive, case by case.

<sup>24</sup> What of vagueness in very simple languages, as in Dorr (2003)’s two word language for declaring how much fruit is on a fruit tree, with a hoot meaning more fruit, a yelp meaning less fruit? Note that I do not require that one’s beliefs be couched in terms of the relevant predicates. A monkey does not have to believe that a fruit tree satisfies the predicate ‘65 pieces of fruit’ to be warranted in exhibiting the  $\Phi(x, \text{‘Yelp’})$  attitudes. It suffices that the monkey believes that the fruit tree has 65 pieces of fruit.



It remains to be seen whether the principle is ultimately tenable. To this end, I will first confront the challenge Weatherston applies to Eklund's **Strong** principle: the possibility of ignorance or error by competent speakers. I will then argue that the principle is extensionally adequate: if we take into account the possibility of ignorance or error, and also control for the role of context, then we can see that the thesis delivers the right results in most ordinary cases. There will of course be some cases, like 'phenomenally conscious', for which we cannot find positive characterizations of borderline cases. But this is a feature of the thesis, rather than a bug: it delivers novel predictions about some hard cases, while also fitting the data concerning ordinary cases.

Whether something is a borderline case of 'water' presumably depends on whether water is a rigid designator or a descriptive term. If 'water' means 'watery' then a suitably viscous mixture of mud and water will be a borderline case. On the other hand if 'water' picks out the actual dominant watery kind, samples consisting of 75 % water and 25 % XYZ may be borderline cases of 'water', despite being perfectly watery.

But reasonable people can disagree about this. Accordingly, reasonable people can disagree about what characterizes a borderline case. But that is perfectly acceptable: I allow that competent speakers may be ignorant or hold false (but reasonable) beliefs about these matters, and accordingly I do not claim that all such speakers are in a position to know which cases are borderline.<sup>25</sup>

When reasonable people disagree about whether 'water' is a rigid designator or a definite description, they disagree (inter alia) about the norms of usage—i.e. about what usage is appropriate. If 'water' rigidly designates the dominant watery kind, then usage norms exist reflecting this fact, and '75 % water-25 % water-look-alike mixture' is a positive characterization of a borderline case of 'water'. This is all that my thesis requires.

As an aside, I stress that my account is compatible with an account based on the dispositions of competent speakers. My notion of a norm of usage could be thought of as equivalent with the notion, employed by Eklund and others, of a usage disposition shared by all fully competent speakers. Here, the distinction between *mere* and *full* competence does much to account for ordinary ignorance and error: insofar as one is ignorant or in error, even if this is reasonable, one is ipso facto less than fully competent. I retain the focus on norms because in the coming sections there is much to say about propositional knowledge *of* the norms, and it is awkward to speak of knowledge of which dispositions are shared by all fully competent speakers. But we may nevertheless think of the standard of dispositions shared by

---

<sup>25</sup> Why doesn't this undermine my claim that competency makes for sensitivity to whether a case is borderline? I take the primary measure of the sensitivity of competent speakers to be whether they comply with relevant norms, not whether they are in a position to know what those norms are. But as I will discuss below, we might think of a fully competent speaker as one who also knows the norms she complies with. In this case, knowing that the norms in question are proprietary of borderline cases (if they are) would put any fully competent speaker in a position to know which cases are borderline.

fully competent speakers as serving to individuate the class of norms we designate as usage norms.<sup>26</sup>

To address another worry about the adequacy of the Positive Characterization Thesis in ordinary cases, we must consider the role of context. ‘5’10” is a positive characterization of a borderline case of ‘tall’ in a context where the comparison class is a roomful of ordinary-sized European adult men, but not where it includes children, or too many professional basketball players. There are various difficulties in saying which cases are borderline that we overcome only by attending to context. For this reason, the Positive Characterization Thesis quantifies over contexts. Now, there are various difficulties in saying just which contextual parameters are relevant to vagueness. And of course, this means that reasonable, competent error and ignorance are possible here too. But I take it that here, too, the facts of the matter are encoded in the relevant usage norms, so this possibility of ignorance and error is not a challenge to the Positive Characterization Thesis.<sup>27</sup>

These considerations suggest that the Positive Characterization Thesis is not too difficult to comply with, and that in ordinary cases of vagueness positive characterizations are not too hard to find.<sup>28</sup> Another worry is that they might be too easy to find: that the thesis turns out to be trivial. I stress that this is not so. If all the thesis required were that there were *something or other* one could learn about  $x$  that implied that ‘ $P(x)$ ’ is borderline, the thesis would be trivial: for “‘ $P(x)$ ’ is borderline’ is something that one might learn about  $x$ . In blocking this sort of objection, the requirement that there exists a suitably *explanatory* usage norm for each borderline case does the heavy lifting. It is because a positive characterization must be explanatory in this sense—the sense in which citing the fact that Al is 5’10” explains why he is a borderline case of ‘tall’—that descriptions that threaten to trivialize the thesis, descriptions that *say* the case is borderline without *showing* it, are not positive characterizations (of which more in Sect. 5).

<sup>26</sup> There is more to be said here. One worry comes from Williamson (2007): Peter the expert logician does not believe (and so does not know) that one ought not infer  $\exists xP(x)$  from  $\forall xP(x)$ . Does he still use the same notions  $\exists$  and  $\forall$  that we use (Boghossian 2010)? If so the link between mastery and norms of usage must be attenuated, or Peter falls short of mastery, despite his expertise (deRosset 2013). Another worry: imagine the reluctant or capricious master, who knows precisely what ought to be done, but has no inclination to do it—an ideally coherent Caligula of grammar. James Joyce and e.e. cummings come to mind. If such a person truly counts as a master, we might do better to say that a norm of usage is a norm that a speaker must know to count as having mastery. This would have the additional benefit of honoring the insight that mastery requires reflective endorsement of one’s practice (as defended in Ginsborg (2012)). I stress again that I do not claim that usage norms explain the facts about meaning and content in general. I allow that in general the norms merely *reflect* semantic facts which are not essentially normative. I make an explanatory claim only in the special case of vagueness.

<sup>27</sup> Brogaard (2010) argues that in: ‘Mary is quite conscious for someone in a vegetative state’, ‘conscious’ is a (contextually sensitive) relative gradable adjective. But as Sebastian (2011) points out, ‘conscious’ does not mean ‘phenomenally conscious’ here.

<sup>28</sup> What of evaluative concepts, like ‘permissible’ or ‘beautiful’ or ‘funny’? Evaluative concepts are very special, and so I do not take a verdict about their vagueness to be incumbent on my defense of the Positive Characterization Thesis. I note if the norms for faultless disagreement were the same as the norms for borderline cases, then the  $\Phi$ -attitudes could not be proprietary of vagueness. But this is best saved for another day.

### 3 The Conceivability Gap is not Merely Semantic

Having defended the Positive Characterization Thesis, I turn now to my defense of the master claim of this paper: that if the Positive Characterization Thesis is true, and if zombies are conceivable, then 'phenomenally conscious' is not vague. The key contention of my defense is that, if zombies are conceivable, then there are too few semantic (or more generally, use-normative) connections between 'phenomenally conscious' and physical, structural or dynamic descriptions for any of the latter to be connected to the former by an explanatory norm of usage. But then no physical, structural or dynamic description<sup>29</sup> can be a positive characterization of a borderline case of 'phenomenally conscious'. I take this contention to be plausible as stated. But to make a more careful case for it, I will unpack the claim that zombies are conceivable as a claim about the limits of what we can know a priori, and then argue that the motivations for accepting this claim are also motivations for accepting a slightly stronger claim.

So let us understand the claim that zombies are conceivable as a claim about a priori epistemology: the claim that the phenomenal truths about the world do not follow a priori from the physical truths.<sup>30</sup> Here, I will focus on a specific consequence of this general claim: that the truth that there is at least one thing that is not determinately a zombie does not follow a priori from the physical truths.<sup>31</sup>

In contrast, claims about usage norms need not be claims about a priori epistemology. Some like Chalmers and Jackson 2001 affirm that semantic principles do follow a priori from the physical truths, supplemented in the right way,<sup>32</sup> but others like Block and Stalnaker (1999) doubt it.<sup>33</sup>

For present purposes we need not resolve this debate. What I do claim is this: we would be unable to deduce facts about consciousness from physical facts even if the usage norms were a priori. We are unable to deduce facts about consciousness from physical facts because there are too few use-normative principles connecting physical descriptions and phenomenal ones, not because there are such principles but they are a posteriori.

Intuitively even masters of the terms involved would be unable to deduce the phenomenal facts from the physical facts.<sup>34</sup> After all, we take ourselves to know how to get on with terms like 'phenomenally conscious', but still are puzzled by the psycho-physical nexus. Further, many experts in the field take it to be a constraint

<sup>29</sup> See Chalmers (2013) for exposition.

<sup>30</sup> Where by 'B follows a priori from A' I mean that the material conditional 'If A then B' is knowable a priori. Cf. Chalmers (2013).

<sup>31</sup> Here, the important claim will be that a positive characterization rules out that its instances are *non-vaguely* P (as opposed to *determinately* P). But this locution is awkward, and many agree that vagueness implies indeterminacy, so I focus here on whether we can rule out that instances are *determinately* P. I do not thereby mean to rule on whether vagueness is incompatible with classical logic.

<sup>32</sup> That is, they follow from the PQTI truths: the physical, phenomenal, totality and indexical truths.

<sup>33</sup> See also Neta 2014 and Schroeter 2014, and Chalmers 2014 for reply

<sup>34</sup> Though see Rabin (2011), Ball (2013), and Alter (2013) for reply.

on an adequate theory of the (meta)semantics of phenomenal concepts, that it explains the existence of the conceivability gap.<sup>35</sup>

‘Phenomenally conscious’ is extraordinary in this respect. Perhaps we do not know a priori whether ‘glassful of a pure sample of the dominant watery kind’ describes a glassful of water, but we have no reason to doubt that there are usage norms which settle the matter.

We are almost there. One final point: in Sect. 2 above, I speculate about whether the  $\Phi(x, 'P')$  attitudes are proprietary of vagueness, in the sense that for any predicate R, if it is appropriate for those who believe that R(x) in context c to  $\Phi(x, 'P')$  in context c, then Rs are borderline cases of ‘P’ in context c. Observe that this is a general claim, making no specific reference to the nature of phenomenal terms of concepts. Accordingly, even if we were given the truth of this proprietary claim, this still would not enable us to deduce the phenomenal facts from the physical, structural or dynamic facts. As we shall see, this will be useful below.

Taken together, what follows from all of this is:

(Deep Gap Principle) For any set of purely physical, structural or dynamic truths, there is no a priori entailment from that set of truths together with a specification of all relevant usage norms, and the supposition that the  $\Phi(x, 'P')$  attitudes are proprietary of vagueness, to the conclusion that not everything is determinately a zombie.

To gloss this in terms of mastery: even masters of all of the terms involved who suppose that the  $\Phi(x, 'P')$  attitudes are proprietary of vagueness cannot rule out that everything is determinately a zombie solely on the basis of any set of purely physical, structural or dynamic truths.

Here my aim in this section has been to present and defend the Deep Gap Principle. I turn now to showing why it entails that there can be no positive characterizations of borderline cases of ‘phenomenally conscious’ in purely physical, structural and dynamic terms.

#### 4 No Positive Characterizations from Across the Gap

While in the previous section I argue that a certain kind of conditional is *not* knowable a priori, in this section I argue that a different kind of conditional is knowable a priori. As I hope to make clear, this latter kind of conditional is a very special case, since it may be transformed into a tautology by the unpacking of terms.

I claim that if ‘R’ is a positive characterization of a borderline case of ‘P’ (in context c), then it follows a priori from the fact that *x* is R in c, together with a

<sup>35</sup> See, e.g., Balog (2012), Papineau (2002). Why is this? Intuitively, it is because of the peculiar sparsity of usage norms for the term. Note that these very considerations have led others, like Papineau (1993, 2002), to conclude that ‘phenomenally conscious’ is indeterminate.

specification of the relevant usage norms and the supposition that the  $\Phi(x, 'P')$  attitudes are proprietary of vagueness, that  $x$  is a borderline case of 'P' in  $c$ .

For example, supposing that the  $\Phi(x, 'P')$  attitudes are proprietary of vagueness, it follows a priori from the fact that Al is 5'10" (together with the usage norms) that Al is a borderline case of 'tall'. Likewise, it follows a priori from the fact that the glass contains a mixture of 75 % water and 25 % another lookalike substance (together with the usage norms, including the fact that 'water' is a rigid designator designating the dominant watery kind) that it is a borderline case of 'glass of water'.

Conditionals of the relevant form amount to tautologies. For as I have defined 'positive characterization of a borderline case', if 'R' is a positive characterization of a borderline case of 'P' (in context  $c$ ), then there is a usage norm stating that it is appropriate for those who believe that  $R(x)$  in context  $c$  to  $\Phi(x, 'P')$  in context  $c$ . But to suppose that the  $\Phi(x, 'P')$  attitudes are proprietary of vagueness is to suppose that for any predicate R, if it is appropriate for those who believe that  $R(x)$  in context  $c$  to  $\Phi(x, 'P')$  in context  $c$ , then Rs are borderline cases of 'P' in context  $c$ . It follows directly that Rs are borderline cases of 'P' in  $c$ . From this, if one also has that  $x$  is R in  $c$ , it follows directly that  $x$  is a borderline case of 'P' in  $c$ .

For this reason I conclude that even those who are generally skeptical of the a priori (but allow that tautologies are a priori) should allow that conditionals of the relevant form are a priori.<sup>36</sup> This is to say that if 'R' is a positive characterization of a borderline case of 'phenomenally conscious' (in  $c$ ) it follows a priori from the fact that  $x$  is R (in  $c$ ) together with a specification of relevant usage norms and the supposition that the  $\Phi(x, 'P')$  attitudes are proprietary of vagueness, that  $x$  is a borderline case of 'phenomenally conscious' (in  $c$ ).

From this, it follows a priori that not everything is determinately a zombie. From this, it follows, given the Deep Gap Principle I defend in the previous section, that no description in purely physical, structural or dynamic language can specify a positive characterization of a borderline case of 'phenomenally conscious'. And note that the claim that the  $\Phi(x, 'P')$  attitudes are proprietary of vagueness merely plays the role of a supposition here: nothing in the argument I have just given depends on its truth.<sup>37</sup>

This rules out almost all of the suggestions that have been made in the literature for how to understand the vagueness of 'phenomenally conscious'. A positive characterization of a borderline case of 'phenomenally conscious' cannot come from a description of a sequence of brains whose neurons are increasingly far apart (staying in touch with nanodevices, perhaps),<sup>38</sup> or from a description of falling

<sup>36</sup> Strictly, a specification of the usage norms could be infinitary, so we appeal to an infinitary conjunction elimination. Otherwise the tautology is classical

<sup>37</sup> If the  $\Phi(x, P)$  attitudes are not proprietary, this means that our account of the constitutive nature of borderline cases is at best only partial. But that is not inconsistent with the Positive Characterization Thesis, nor with the conclusions I will draw below. Note also that the conclusion of this section does not depend on the truth of the Positive Characterization Thesis, i.e. the thesis that every borderline case has a positive characterization. In this section I exploit my definition of a positive characterization, but do not rely on the claim that every borderline case has one.

<sup>38</sup> Zuboff (1981). See also Bostrom (2006).

asleep (in physical or functional terms), or from a description of a signal on the cusp of being discriminable from noise by subpersonal perceptual systems, or of a message that is on the cusp of being broadcast to enough systems to count as being broadcast globally.<sup>39</sup> A positive characterization of a borderline case of ‘phenomenally conscious’ cannot come from any sequence of descriptions of persons in increasingly severe minimally conscious or vegetative states (provided that those descriptions are in physical and functional, rather than phenomenal terms), and it cannot come from any analysis of the cusp of our ordinary behavioral and functional criteria for ascribing consciousness, or of creatures who fit there, such as fish or slugs. No description of the Freudian subconscious (understood in a functional way) will do. Nor will any description of a sequence of systems of decreasing functional complexity, or informational complexity (functionally construed), not even if the description covers all of the facts about the world in question (in materialistic terms).<sup>40</sup>

Why doesn’t this prove too much? The Deep Gap principle exploits a very special feature of the case of ‘phenomenally conscious’: we have independent reason for thinking that the collection of usage norms connecting physical, structural and dynamic terms with phenomenal ones is very thin on the ground. We have no reason for suspecting anything analogous in most cases.

Moreover, even in cases where we do accept some analogue of the Deep Gap principle, and so, accept that positive characterizations of borderline cases of a predicate must enjoy some kind of a priori connection to that predicate, still this generally leaves many options open. For example, even if colors are primitive, still we may positively characterize a borderline case of ‘orange’ by ‘rgb(255,69,0)’ (a shade on the borderline between orange and red), or ‘*this shade*’ (a perceptual demonstrative for rgb(255,69,0)). Indeed, even for the case of ‘phenomenally conscious’ itself there is still a range of candidate positive characterizations of borderline cases that the test I have presented in this section does not rule out. I now turn my attention to these.

## 5 No Positive Characterizations from the Phenomenal Side of the Gap

In the previous sections we have seen that if a description is purely physical, structural and dynamic, then it cannot be a positive characterization of a borderline case of ‘phenomenally conscious’. But might some other description carve out space for borderline cases?

Here I’ll argue that the leading candidates fall short. I will consider three classes. First, descriptions that imply that *x* is a borderline case, but by saying rather than

<sup>39</sup> Brogaard (2010).

<sup>40</sup> Finally, note that this holds even if we believe it is a priori that a system of low enough functional complexity is not phenomenally conscious, since even then, a description of a system of whatever complexity level might still be a description of perfect zombies. Cf Antony (2006a, 2008) whose argument that PHENOMENALLY CONSCIOUS is not vague hinges on that concept containing “no material elements.”

showing, and so are not positive characterizations of borderline cases. Second, descriptions that pass the buck, offering what might be positive characterizations, but only by transferring the problem of finding positive characterizations elsewhere. Third, descriptions that exploit our ignorance about where borderline cases might lie.

Of the first sort are descriptions like 'borderline case of 'phenomenally conscious''. Obviously this description's instances are borderline cases, but just as obviously, the description is not a positive characterization, because of the explanatory constraint on positive characterizations. Recall that where 'R' is a positive characterization of a borderline case of 'P', the fact that it is appropriate to adopt the  $\Phi(x, 'P')$ -attitudes toward things one believes to be R *constitutively explains* why satisfying 'R' secures that something is a borderline case of 'P'. But disquotational principles suffice to explain why satisfying 'borderline case of 'P'' secures that something is a borderline case of 'P'.

The same reasoning covers more sophisticated examples. Consider the description ['borderline case of 'roughly 40Hz oscillation', and such that roughly 40Hz oscillation = phenomenal consciousness']. Even supposing that this description is satisfiable, the existence of a norm of usage recommending adopting the  $\Phi(x, 'phenomenally conscious')$ -attitudes toward things one believes to be [borderline cases of 'roughly 40 Hz oscillation' and also such that roughly 40 Hz oscillation = phenomenal consciousness], does not play any role in explaining why satisfying this description secures that something is a borderline case of 'phenomenally conscious'. Rather, we have all the explanation we need in terms of the logic of the description itself, since we may directly derive that anything that satisfies it is a borderline case of 'phenomenally conscious', assuming the right principles for reasoning with vagueness or indeterminacy in identity claims.<sup>41</sup> Matters would be different if the identity were a priori or a consequence of usage norms. For then we might argue that somehow usage norms governing positive characterizations of borderline cases of '40 Hz oscillation' explain why satisfying those positive characterizations suffice to secure that x is a borderline case of 'phenomenally conscious'. But no such identity is likely to be a priori or a consequence of usage norms.

Of the second sort are descriptions that exploit the relationships between the term 'phenomenally conscious' and other terms that are or may be semantically connected to it, like for example 'feeling pain' or 'has intentional states'. E.g., it might be an explanation of what the vagueness of a case of 'phenomenally conscious' consists in, that it satisfies some description like 'x definitely either feels pain or nothing at all, but it is vague whether x is feeling pain'.<sup>42</sup>

Similarly, it might be that there are semantic connections between 'phenomenally conscious' and descriptions of some kind of intentional state  $\psi$  such that it is settled

<sup>41</sup> The question is whether the identity is to be read as asserting ' $\text{Phenomenally Conscious}(x) \Leftrightarrow \text{Roughly 40Hz Oscillation}(x)$ ', or as asserting ' $\Delta\text{Phenomenally Conscious}(x) \Leftrightarrow \Delta\text{Roughly 40Hz Oscillation}(x)$ '. If the latter, then satisfying the description may not even entail that x is a borderline case of 'phenomenally conscious'.

<sup>42</sup> See Sebastian (2011).

by the physical facts plus the usage norms that one has  $\psi$ -intentional states if and only if one is phenomenally conscious. If so then ‘borderline case of having  $\psi$ -intentional states’ might be a positive characterization of a borderline case of ‘phenomenally conscious.’

But if so, we face new problems as difficult as the original one: that of specifying a positive characterization of a borderline case of ‘feels pain’ that is compatible with ‘definitely feels pain or nothing at all’, or specifying a positive characterization of a borderline case of ‘has  $\psi$ -intentional states’.

Of the third sort are descriptions which tempt us to think that vagueness may be afoot without settling the matter. Examples include descriptions of gradually falling asleep, or of being in the related state Italians call *dormiveglia*. Tye (1996, 2000) suggests a case where you hear quieter and quieter sounds (through a headphone, as a subject in a psychological experiment) until you reach a point where you are unsure if you have heard a sound or not. Along similar lines are cases involving the hard to quantify threshold between conscious and unconscious perception.

Another source of such cases is philosophical dispute. For example, can there be consciousness without attention or without access? Can there be qualitative experience that does not impart a single subjective perspective? Can there be purely cognitive phenomenology? Can there be entirely instantaneous experiences, or must experience take time? You might think that at least some of these questions appear to be intractable because ultimately there is no fact of the matter.

In reply, I say that, while some of these cases may well be borderline cases of ‘phenomenally conscious’, we lack any reason for thinking of the descriptions I have indicated above as positive characterizations. To the contrary, each of these descriptions appears to be fully consistent with there being determinate cut-offs of which we are ignorant (see the related discussion in Antony (2006a, 2008)).

Where then do we stand? There are regrettably no general principles for ruling out positive characterizations that involve more than purely physical, structural or dynamic terms. But I have attempted to classify those candidates that anyone in the literature has come close to articulating, and show why they all are wanting.

This leaves one salient possibility open: that positive characterizations of borderline cases of ‘phenomenally conscious’ exist but must be couched in vocabulary that we have not yet developed. I take this possibility seriously, because I allow for the possibility that the norms of usage for the term ‘phenomenally conscious’ involve hidden features that we do not, at present, fully understand.<sup>43</sup>

There is a school of thought, Russellian monism,<sup>44</sup> according to which we will only understand the emergence of consciousness if we come to grasp the intrinsic natures of the physical properties that constitute it. According to this line of thought,

<sup>43</sup> Compare Antony (2006a, 2008)’s discussion of LIFE as deployed in 1750. In 1750, people were not disposed to recognize viruses as borderline cases, but now we are. Antony considers this a change of concept. But Antony does not adequately distinguish between the possibility that one day we will use the words ‘phenomenally conscious’ to express a concept with different usage norms which is vague, and on the other hand the possibility that we will one day notice a way in which the term we have been using all along has been vague, unbeknownst to us. The latter is my concern.

<sup>44</sup> See for example Pereboom (2011), Stoljar (2001).



there is more to the physical than its structure and dynamics, and this extra intrinsic component is what our explanation of consciousness requires. Note that there are both metaphysical and epistemic-semantic components to this picture: it calls for a metaphysical contrast between structure and dynamics on the one hand and intrinsic natures on the other, and it calls also for the concepts of these latter to give us a richer, more nuanced understanding of 'phenomenally conscious'. If both of these claims hold, then positive characterizations of borderline cases might turn out to be formulable in terms of these new concepts of the intrinsic natures of things. This would mean that 'phenomenally conscious' does have borderline cases after all, even though we today lack the concepts to appreciate why.

However, these claims which underwrite Russellian monism are very speculative, and for this reason I am content to allow that the truth of Russellian monism (comprising both its metaphysical and its semantic/epistemic theses) might well lead to an exception to my master claim that 'phenomenally conscious' has no borderline cases. That is, I am content to have made the case here that 'phenomenally conscious' has no borderline cases *if* ordinary materialism (or non-Russellian non-materialism) is true.

## 6 Conclusion

I have argued that 'phenomenally conscious' is not vague. But this leaves open that it is indeterminate in some other way. There are a number of ways for sentences to be indeterminate, or to lack a truth-value, without being vague, arguably including presupposition failure, the open future, quantum indeterminacy, various forms of relativism or context-sensitivity, and finally the radical sort of metasemantic indeterminacy considered by Kripke, Quine and others. But the claims about 'phenomenally conscious' that I am interested in do not appear to involve presupposition, we may restrict our attention to those in the present tense, and there is no obvious role for any sort of relativism or context-sensitivity here, beyond those which we have discussed above. I leave the more subtle discussions of quantum indeterminacy and radical metasemantic indeterminacy for another day.

**Acknowledgments** Acknowledgments to audiences in Canberra, New York, Sydney, Toronto, Wagga Wagga, and Wollongong, and in particular to Eli Alshanetsky, Michael Antony, Ned Block, Paul Boghossian, David Chalmers, Kit Fine, Franz Huber, Geoffrey Lee, Dan Lopez de Sa, Colin Marshall, John Morrison, Diana Raffman, Michael Raven, Stephen Schiffer, Miguel Sebastian, Ted Sider, Nicholas J.J. Smith, Sharon Street, Peter Unger, Robbie Williams and Achille Varzi.

## References

- Alter, T. (2013). Social externalism and the knowledge argument. *Mind*, 122(486), 481–496.
- Alston, W. (1964). *Philosophy of language*. Englewood Cliffs, NJ: Prentice-Hall.
- Antony, M. (2008). Are our concepts conscious state and conscious creature vague? *Synthese*, 68(2), 239–263.
- Antony, M. (2006a). Vagueness and the metaphysics of consciousness. *Philosophical Studies*, 128, 515–538.
- Antony, M. (2006b). Papineau on the vagueness of phenomenal concepts. *Dialectica*, 60(4), 475–483.

- Ball, D. (2013). Consciousness and conceptual mastery. *Mind*, 122(486), 497–508.
- Balog, K. (ms.). Hard, harder, hardest. <http://philpapers.org/rec/BALHHH>. Accessed 23 Sept 2016.
- Balog, K. (2012). In defense of the phenomenal concept strategy. *Philosophy and Phenomenological Research*, 84(1), 1–23.
- Block, N. (2002). Some concepts of consciousness. In D. Chalmers (Ed.), *Philosophy of mind: Classical and contemporary readings*. Oxford: Oxford University Press.
- Block, N., & Stalnaker, R. (1999). Conceptual analysis, dualism, and the explanatory gap. *The Philosophical Review*, 108, 1–46.
- Bostrom, N. (2006). Quantity of experience: Brain duplication and degrees of consciousness. *Minds and Machines*, 16, 185–200.
- Boghossian, P. (2010). A companion to epistemology. In Bernecker & Pritchard (Eds.), *Williamson on the A Priori and the Analytic*. London: Routledge.
- Boghossian, P. (2003a) Epistemic analyticity: A defense. *Grazer Philosophische Studien*, 66, 15–35.
- Boghossian, P. (2003b). The normativity of content. *Philosophical Issues*, 13, 31–45.
- Brandom, R. (1994). *Making it explicit*. Cambridge, MA: Harvard University Press.
- Brogaard, B. (2010). Degrees of consciousness. In *Presented at the SPAWN conference on metaphysics*, Syracuse University. Accessed at <https://sites.google.com/site/brogaardb/home>.
- Chalmers, D. (2013). *Constructing the world*. Oxford: Oxford University Press.
- Chalmers, D., & Jackson, F. (2001). Conceptual analysis and reductive explanation. *Philosophical Review*, 110, 315–361.
- Dennett, D. (1998). Animal consciousness: What matters and why. In D. Dennett (Ed.), *Brainchildren*. Cambridge: MIT Press.
- deRosset, L. (2013). *Analyticity and ontology*. OUP: Oxford Studies in Metaphysics.
- Dummett, M. (1975). Wang's paradox. *Synthese*, 30, 301–324.
- Dorr, C. (2003). Vagueness without ignorance. *Philosophical Perspectives*, 17(1), 83–113.
- Eklund, M. (2005). What vagueness consists in. *Philosophical Studies*, 125, 27–60.
- Fara, D. G. (2000). Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics*, 28, 45–81.
- Fine, K. (1975). Vagueness, truth and logic. *Synthese*, 30, 265–300.
- Gaifman, H. (2010). Vagueness, tolerance and contextual logic. *Synthese*, 174, 5–46.
- Ginsborg, H. (2012). The normativity of meaning. *Proceedings of the Aristotelian Society (Supplementary Volume)*, 86, 127–146.
- Glüer, K., & Wikforss, Å., (2015). The normativity of meaning and content. The Stanford Encyclopedia of Philosophy, In Edward N. Zalta (Ed.), <http://plato.stanford.edu/archives/sum2015/entries/meaning-normativity/>.
- Goff, P. (2013) Orthodox property dualism + linguistic theory of vagueness = panpsychism. In Brown, R. (Ed.) *Consciousness inside and out*. Springer.
- Greenough, P. (2003) Vagueness: A minimal theory. *Mind*, 112(446).
- Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30, 1–45.
- Kripke, S. A. (1980). *Naming and necessity*. Cambridge, MA: Harvard University Press.
- Kripke, S. A. (1982). *Wittgenstein on rules and private language*. Cambridge, MA: Harvard University Press.
- Lewis, D. (1975). Languages and language. In K. Gunderson (Ed.), *Minnesota studies in the philosophy of science* vol. VII, (pp. 3–35). Minneapolis: University of Minnesota Press.
- McGinn, C. (1996). *The character of mind* (2nd ed.). Oxford: Oxford University Press.
- Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83, 435–450.
- Papineau, D. (2002). *Thinking about consciousness*. Oxford: Oxford University Press.
- Papineau, D. (1993). *Philosophical naturalism*. Oxford: Blackwell.
- Pereboom, D. (2011). *Consciousness and the prospects of physicalism*. New York: Oxford University Press.
- Rabin, G. (2011). Conceptual mastery and the knowledge argument. *Philosophical Studies*, 154, 125–147.
- Raffman, D. (1996). Vagueness and context-relativity. *Philosophical Studies*, 81(2–3), 175–192.
- Schiffer, S. (2003). *The things we mean*. Oxford: OUP.
- Shapiro, S. (2006/2008). *Vagueness in context*. Oxford University Press.
- Searle, J. (1992). *The rediscovery of the mind*. Cambridge, MA: MIT Press.
- Sebastian, M. (2011). Self involving representationalism: A naturalistic theory of phenomenal consciousness. Doctoral Dissertation, LOGOS, University of Barcelona.

- Smith, N. J. J. (2005). Vagueness as closeness. *Australasian Journal of Philosophy*, 83(2), 157–183.
- Stoljar, D. (2001). Two conceptions of the physical. *Philosophy and Phenomenological Research*, 62, 253–281.
- Tye, M. (2000). *Consciousness, color, and content*. Cambridge, MA.: MIT Press.
- Tye, M. (1996). Is consciousness vague or arbitrary? *Philosophy and Phenomenological Research*, 56(3), 679–685.
- Unger, P. (2004). The mental problems of the many. In D. Zimmerman (Ed.), *Oxford studies in metaphysics* (Vol. 1). Oxford: Clarendon Press.
- Unger, P. (1988). Conscious beings in a gradual world. *Midwest Studies in Philosophy*, 12(1), 287–333.
- Weatherson, B. (2010). Vagueness as indeterminacy. In Dietz & Morruzzi (Eds.), *Cuts and clouds*. Oxford: Oxford University Press.
- Weatherson, B. (2003). Many many problems. *Philosophical Quarterly*, 53, 481–501.
- Williamson, T. (2007). *The philosophy of philosophy*. New York: Blackwell.
- Williamson, T. (1994). *Vagueness*. London: Routledge.
- Wright, C. (2001). On being in a quandary: Relativism, vagueness, logical revisionism. *Mind*, 110, 45–98.
- Wright, C. (1975). On the coherence of vague predicates. *Synthese*, 30, 325–365.
- Zuboff, A. (1981). The story of a brain. In D. R. Hofstadter & D. C. Dennett (Eds.), *The Mind's I. Basic Books* (pp. 202–212).