Philosophy Graduate Theses & Dissertations

Philosophy

Spring 1-1-2014

Vagueness and Fundamentality

Joseph Michael Fraley
University of Colorado at Boulder, joefraley08@hotmail.com

VAGUENESS AND FUNDAMENTALITY

Joe Fraley

B.A., Western Washington University, 2011

M.A., University of Colorado, Boulder, 2014

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirement for the degree of

Master of Arts

Department of Philosophy

2014

This thesis entitled:

Vagueness and Fundamentality

written by Joseph Michael Fraley Jr.

has been approved for the Department of Philosophy

David Barnett

Graham Oddie

Michael Tooley

Date_____

The final copy of this thesis has been examined by the signatories, and we

Find that both the content and the form meet acceptable presentation standards

Of scholarly work in the above mentioned discipline.

Fraley, Joseph Michael (M.A., Philosophy)

Vagueness and Fundamentality

Thesis directed by Professor David Barnett

ABSTRACT

How can we tell whether a concept is vague? Is it possible to generate reliable methods for demonstrating vagueness and clearness without having an analysis of the nature of vagueness? Are there good methods that are neutral with respect to standard metaphysics of vagueness? I argue for two methods. First, I argue that the way sorites series work entails that no concepts that pick out fundamental features are vague. If a concept is vague, then what it picks out is not fundamental. Second, I argue that if a concept is not vague, then what it picks out is fundamental. There are wide implications of this view for our prospects of success elsewhere in metaphysics, metaethics, and the philosophy of mind.

ACKNOWLEDGMENTS

CONTENTS

"I warn the reader that this chapter has to be read carefully, and that I lack the art of being clear to those who are not willing to be attentive." – Jean Jacque Rousseau

In this essay, I consider one main question. How can we tell whether a concept is vague? I will explore this question *without* committing to any particular analysis of the nature of vagueness. The only features of vagueness that I will assume are those that are perfectly general to any theory. This question makes up the first half of the essay. In the second half, I ask how my answers to this question shed light on other philosophic questions.

I'll make one important restriction on my claims about vagueness. The kind of vagueness I have in mind is the kind that is traditionally demonstrated using a sorites series.[1] The kind of vagueness I have in mind is demonstrated in the following way. The concept of **being bald**[2] clearly applies to a man with no hairs. It clearly applies to a man with two hairs, and three hairs, and so on. At a certain point in this series, the concept of **being bald** clearly does not apply. For the range in between clear cases of **being bald** and clear cases of **being not bald**, it is vague whether he's bald. Again, this use of vagueness is intended to be neutral between different readings, such as "it is indeterminate whether he's bald," or "it is impossible to know whether he's bald," and so on. I take it that all of our ordinary vague concepts are

---

[1] This use of vagueness is explicitly preferred by Shafer-Landau in his (1994 and 1995). It is not explicitly stated, but I take it that it is also the use preferred by Antony (2006), Barnett (2009), Sainsbury (1996), Sorensen (1985), Tye (1994), and others.

[2] In order to distinguish that I am referring to a concept, rather than a feature picked out by that concept, or a word that employs the concept, I will use boldfaced type and quotation marks. For example, I use the words "being bald" to talk about the concept of **being bald**, which picks out the feature of being bald. I borrow this convention from Michael V. Antony (2008).

susceptible to these demonstrations – for example, the concept of **being bald**, the concept of **being tall**, or the concept of **being young**. If there are other non-sorites kinds of vagueness, then my arguments will not concern them. In my writing I will use the term "clear" as follows: a clear concept is one that is not at all vague. The less clear something is, the more vague it is. For an object to be clearly *F* is for the object to be *F,* and for there to be no vagueness about its being *F*.

§1 – The Main Argument

Consider the concept of **being bald**. This concept clearly applies to some things, clearly fails to apply to some things, and it's not clear whether it applies to others – typically referred to as "borderline" cases. In other words, there is vagueness as to whether the concept applies to these so-called borderline cases. If a man has no hairs on his head, then he is clearly bald. If his head is covered with hair, then he is clearly not bald. But there are lots of men about whom we're uncomfortable saying whether they're bald because it seems to us unclear, or vague, whether they are bald. Contrast this with concepts that people have argued are *not* vague. For example, it's been argued that the concept of **existing** is not vague, that the concept of **personal identity** is not vague, and that the concept of **moral obligation** is not vague.[3] If we think that some concepts are vague and some aren't, then there must be *something* cueing us into the difference. Why does a sorites series reveal that baldness is vague, but not that

---

[3] Ted Sider (1997), Ryan Wasserman (2012), and Thomas Reid (1785) respectively.

personal identity is vague? What kinds of questions can we ask about concepts that will help us

decide which are which?

One feature of vague concepts is that they admit sorites series. A sorites series is a list

of cases $\langle a_1 - a_n \rangle$, where a_1 is a case in which it's clear that a particular concept applies, a_n is a

case in which it clearly doesn't, and for some range in the middle it isn't clear either way. In a

sorites series about baldness, we'd say that in a_1 the man has no hairs; he is clearly bald. In a_2,

the man has two hairs. In a_3, the man has three hairs, and so on. We might end the series with

a case in which a man has a million hairs, making him clearly not bald. Somewhere in the

middle of the series is a range of cases in which it is not clear that he is bald and not clear that

he is not bald. That is to say, it is vague whether he is bald.

In the above sorites series, I have described some conditions under which the concept

of **being bald** clearly applies, and I have described some conditions under which the concept

clearly does not apply. Whether the concept of **being bald** applies to a man depends on the

number of hairs on his head (and perhaps also on how the hairs are arranged, but for simplicity

I will ignore this factor). The concept of **being bald** clearly applies to the man in the first case,

because he has no hairs. The concept of **being bald** clearly does *not* apply in the last case,

because the man has a great many hairs. It's not clear whether the concept applies to men in

the middle cases. Thus, the description of the application conditions for the concept of **being

bald** are not clear. Constructing a sorites series requires that one is able to give a description of

a concept's application conditions that is unclear in the above fashion.

One cannot construct a sorites series to demonstrate that a concept is vague without at least *some* way of describing its application conditions that doesn't use the concept itself in the description. For example, consider a concept that appears to be perfectly clear: the concept of **existing**.[4] In order to construct a sorites series, we require a series of cases $\langle a_1\text{-}a_n \rangle$ where it is clear at each end whether the concept applies, but not clear for a range in the middle. But it's not enough to merely say that in the first case something clearly exists, in the second case it's vague whether something exists, and in the final case clearly nothing exists. That series doesn't do anything to *show* that the concept of **existing** is vague. It doesn't give us any sense of what changes in the second case that makes it vague whether the concept applies. We need to hear that some conditions change from case to case in order to make an assessment about whether the concept applies. The series above just stipulates that the concept of **existing** is vague.

Of course, providing application conditions for the concept of **existing** is at the very least extremely hard and plausibly impossible. Indeed, it's not obvious where one would even begin – besides trivial statements like "an object exists only if it is real," where something is real only if it exists. This important rule for building a sorites series is the key to my suggestion for how to distinguish vague concepts from clear ones. Demonstrating a sorites series using the concept of **existing** would require a non-trivial description of its application conditions – describable in terms other than "exist" or its cognates. Being able to demonstrate such a series

---

[4] Throughout this work, I will assume that there is a genuine concept of **existing**. It has many of the features that other conceptual predicates have. For example, existence can be naturally transformed into a gerund: as in the concept of **existing**. Nothing turns on this. Its purpose here is merely illustrative, but its every instance can be substituted for whatever the reader takes to be a clear concept – such as the concept of **having charge**, or the concept of **being a fundamental particle**.

would require, for example, that whether something exists depends on how many hairs it has. On such a description, nothing can exist with zero hairs, and the more hairs you have the more clearly you exist. That there don't seem to be any such application conditions is exactly what's in the way of constructing the sorites series about **existing**. Without such a description, there is no non-question-begging way of distinguishing cases in a series with respect to their existing. It's our description of **being bald** in terms of **having so many hairs** that elicits the certainty that some points in the series must be special. Without this description, we would have no way of predicting that a certain range of the series would make us uncomfortable about clearly saying whether **being bald** applied in those cases.

Many concepts pick out features of objects or the world. For example, the concept of **being bald** picks out a feature of humans – it picks out a way their heads might be. For many concepts, the ways we describe their application conditions also name features of objects or of the world. For example, the concept of **being bald** only applies to things with a certain number of hairs on their head. Our ordinary description of non-trivial application conditions for baldness names features of humans – it names their number and arrangement of hairs.

Thus, for concepts like **being bald**, we find a special kind of dependence between the features they pick out and the features named when we describe their application conditions. Whether or not an object has the feature picked out by the concept *depends* on whether or not an object has the needed features named in the description of the concept's application conditions. If the application conditions were trivial or circular, then this dependence would also be trivial. For example, we apply the concept of existence to something only if it has the

feature of being real, but this is an empty, uninteresting kind of dependence. In contrast, we

apply the concept of baldness to something only if it has features like having very few hairs and

having most of its scalp exposed. This kind of dependence is more substantive. Whether

something is bald depends trivially on whether it's bald, but it also depends substantively on

other details involving its specific hair situation.

For any feature, if it only obtains as a result of other features in this way, then the

feature is *not* fundamental. For my purposes, it's enough that fundamentality is a metaphysical

dependence *of some kind*. For example, tables are probably not fundamental entities. Tables

depend for their existence on the existence and arrangement of their materials. But no amount

of wood – for example – is a fundamental entity either, since wood depends for its existence on

the existence of certain molecular structures. But no molecular structure is fundamental, since

molecular structures depend for their existence on the existence and arrangement of certain

atoms. If there's a point at which this series ends, it's at a fundamental entry. That is, if

anything on the list did not depend for its existence on anything else, that would be the end.

This is the kind of dependence relation I mean when I claim that something is not fundamental.

If there are no such things, then for my purposes nothing is fundamental.[5]

In the same way, we see that the concept of **being bald** does not pick out a fundamental

feature. Baldness is not a fundamental feature, since baldness *depends* for its presence on the

presence of other features, like number and arrangement of hairs. The concept of **being a hair**

---

[5] For a representative list of features that philosophers have generally considered fundamental,
see Alexis Burgess (2012, 228). It should be clear that I have absolute fundamentality in mind,
although the relative notion of fundamentality can be illustrative as well.

does not pick out a fundamental feature either – for the same reason. The feature of being a hair depends for its presence of other features, like the presence of certain molecular structures. And so on.

To summarize – the vague concepts can be demonstrated using a sorites series. One can only construct sorites series if one can give a non-trivial description of the concept's application conditions – that is, if one can describe application conditions for the concept without using the concept itself. When the application conditions of a concept admit descriptions of this kind, there is a special kind of dependence between the feature the concept picks out and the features specified by our description of the concept's application conditions. This dependence is the kind relevant to fundamentality. Nothing is fundamental if it depends in this way on the presence of other features. Thus, if a concept is vague, then it is *not* fundamental. So, if a concept picks out something that's fundamental, then the concept is *not* vague.[6] Here, then, is my main argument:

The Main Argument

> P1. If a concept is vague, then its vagueness can be demonstrated using a sorites series.
>
> P2. If a concept's vagueness can be demonstrated using a sorites series, then its application conditions can be described without using the concept itself.

---

[6] This claim is just the contraposition of the Main Argument's conclusion. You could prove that this conclusion follows just by inverting the Main Argument's direction. Since I started by talking about vague concepts rather than fundamental features though, I have presented the argument in this direction for the sake of continuity.

P3. If a concept's application conditions can be described without using the concept itself, then the features that determine whether the concept applies are more fundamental than what the concept picks out.

P4. If the features that determine whether the concept applies are more fundamental than what the concept picks out, then whatever the concept picks out isn't fundamental.

P5. Thus, if a concept is vague, then whatever it picks out is not fundamental.

My conclusion seems independently plausible when we consider the baldness series. Whether someone is bald depends on how many hairs are on their head. Thus, given what I said above, being bald is a less fundamental feature than is having a hundred hairs. It seems that it's less vague how many hairs you have than it is vague whether you're bald. I will say more about the relative vagueness of concepts in §3, but for now I think this statement is intuitive enough. It's not that in every case it's less vague how many hairs one has than it is whether one is bald. It's more like the concept of **being bald** has lots of borderline cases independently of borderline cases of **being a hair**. For example, when the number of hairs on a person's head is clearly 5,000, it's vague whether they are bald. This is a case where it's clear how many hairs the person has, but not whether the man is bald. But *on top* of all such borderline cases, the concept of **being bald** has borderline cases that are a result of it being vague how many hairs as a man has. In the same fashion, whether you have a hundred hairs depends on whether there are a billion molecules arranged hair-wise on your head. Having a

hundred hairs is less fundamental than having a billion molecules arranged as hairs. And again,

number and arrangement of molecules are less vague than number and arrangement of hairs.

§2 – Potential Problems

I will now consider possible objections to The Main Argument. The argument is a series

of valid hypothetical syllogisms, so there are no objections to its validity. I'll begin by evaluating

each premise. (P1) says that if a concept is vague, then its vagueness can be demonstrated by a

sorites series. A counter example to (P1) would be a vague concept whose vagueness cannot

be demonstrated using a sorites series, which I have explicitly said I am not considering. This is

my assumed starting point; it limits what kind of vagueness I address.

(P2) says that if one can demonstrate a concept's vagueness using a sorites

series, then its application conditions can be correctly described without using the

concept itself. In other words, the concept must permit non-trivial descriptions of its

application conditions.   In order to produce a counter example to (P2), one must

furnish a sorites series for a concept whose application conditions only permit trivial

descriptions. One candidate is the concept of **being vague**. The concept of **being vague**

is plausibly vague itself, even though it isn't obvious that its application conditions

permit non-trivial description.

The sorites series for the concept of **being vague** is a second-order sorites series,

in that it shows how *vagueness* is vague. The second-order series can be constructed

from any other first-order sorites series. For example, consider the baldness series. In

the first case the man is clearly bald, in the last case the man is clearly not bald, and in

the center-most case it's not clear whether he's bald. In some case in between – let's

say a_{5000} – it is *clearly* vague whether the man is bald. Suppose that 75% of his scalp is

showing through sparse hair. The concept of **being vague** *clearly* applies in this case.

But if we take a step in either direction, it gets less vague whether the man is bald.   In

a_{5001} the man has 74.8% of his scalp showing through hair a little less sparse. It is still

clearly vague whether the man is bald, but a little less so. In a_{5002} the man has 74.6% of

his scalp showing through hair a little less sparse, and so on.

A useful way of thinking about the sorites series for the concept of **being vague**

is that it begins with the member of the baldness series that is most vague. It is least

clear whether the concept of **being bald** applies in a_{5000}. However, a_{5000} is the case in

which the concept of vagueness *most* clearly applies, so I'll refer to it as b_1 in the

second-order series. In a_{5001} of the baldness series, it's a little more clear whether the

concept of **being bald** applies. That means that it's a little *less* clear whether the

concept of **being vague** applies. Call this case b_2 in the vagueness series. The last

member of the vagueness series is a clear member of the baldness series. In a_{9999} of the

baldness series – b_{4999} in the second-order series – the man's head is covered in hair. He

is clearly not bald; it is clearly not vague.

Here is the worry for (P2). Since the vagueness series is just the baldness series

renumbered, the vagueness series does not require describable application conditions for the

concept of vagueness. The baldness series can be used to demonstrate that vagueness is a vague concept, and for that you only need to describe application conditions for the concept of baldness. In response, I must either show that one cannot use a sorites series to demonstrate that the concept of **being vague** is vague, or that you really *do* need a non-trivial description of the application conditions of **being vague** in order to generate the second-order series.

The first response is philosophically unsatisfying. The second-order series is structurally analogous to the first-order series, which provides some evidence that it works the same way. In addition, we get the same discomfort naming clear borders in both series. Drawing a clear line between the last case of baldness and the first case of non-baldness feels wrong. So, the concept of baldness must be vague. This discomfort is how the series demonstrates the concept's vagueness. So too, drawing a clear line between the last clear case in the series and the first vague case feels wrong. The concept of vagueness must itself be vague. So the two series have structural similarities, and they produce similar effects in us. It's unlikely that the concept of vagueness cannot be used in a sorites series.

You might say that the second-order series can't succeed because the concept of vagueness isn't really vague. In other words, there is no second-order vagueness. This response amounts to simply biting the bullet, since I have just tried to show that the second-order series makes us uncomfortable in the same manner as the first-order series.[7]

---

[7] Robert Deas is at least attracted to this view: "We have no difficulty in recognizing or imagining border- line cases of 'small', but the idea of something which is neither clearly vague nor clearly otherwise is incomprehensible…" (1989, 29).

[11]

However, the second response is more promising than the first, and it's the one I endorse. Although it sounds at first like baldness is the only concept being described in the second-order series, really vagueness is being described incognito. Here's what I mean. The first-order series is about baldness, and each member describes a case using the conditions under which baldness applies. Each member says something about the number and arrangement of hairs on each man's head. For example, a_1 tells how many hairs the man has and how they're arranged, rather than just saying "he's bald." The second-order series is about vagueness, and each member describes a case using the conditions under which vagueness applies. However, the descriptions still seem to refer to the number and arrangement of hairs on each man's head. b_1 tells us that the 75% of the man's scalp is visible through sparse hair. But surely the number and arrangement of a man's hairs are not a part of the application conditions for vagueness. After all, vagueness applies to all sorts of concepts besides **being bald** – such as **being rich**, and **being tall**.

What's really being described in the second-order series is our discomfort with saying whether it's vague whether the man is bald. The series requires cases of bald men in order to elicit a verdict about whether it's vague whether they're bald. The descriptions of the man's head in each case only serve to score how comfortable we are with saying that the man is bald. My suggestion is that the discomfort being tracked *is* a plausible part of application conditions for vagueness. It is vague whether *X* only if competent speakers are uncomfortable asserting *X*.

Here's another way of making the same point. When we hear the first-order series, we can ask in what sense it tells us anything about baldness. How are the descriptions of each case in the series related to baldness? Why does the series tell us that the concept of baldness is vague, rather than some other concept? The answer is that the cases describe part of *what it is* to be bald. Having a certain number of hairs is just a part of *being* bald (maybe all there is to being bald).

Similarly, when we hear the second-order series, we can ask in what sense it tells us anything about vagueness. How are the descriptions of each case in the series related to vagueness? Why does the series tell us that the concept of vagueness is vague, rather than some other concept? In particular, why does it tell us something about vagueness rather than baldness? The answer is that the conditions being described in each case are a part of *what it is* to be vague. The series is describing cases in terms of the confidence with which competent English speakers ascribe vagueness to each case. Of course, that might not be *all there is* to vagueness. The confidence of language users might not be a necessary and sufficient condition for the concept of vagueness, but it is surely a big part of vagueness.

The second-order series is building this condition into the descriptions of each case. For each member of the second-order series – although it only *sounds* like number and arrangement of hairs are being described – in reality this is just a convenient way of getting at how comfortable competent speakers are about saying whether the man is clearly bald. Thus, it's false that you can describe a sorites series

about the concept of vagueness without *somehow* describing how vagueness works. Namely, you need to describe cases that determine the comfort of competent speakers with ascribing some vague concept.

(P3) says that if you can describe the concept's application conditions non-trivially, then the features that determine whether the concept applies are more fundamental than what the concept picks out. Obviously, all concepts apply to something only if their application conditions are satisfied. The point is that their application *depends* on the world matching the way the conditions are described. The application of baldness depends on the number and arrangement of hairs on a human's head. This means that the feature picked out by the concept of baldness is a less fundamental feature than those picked out by number, arrangement, and hair.

Two challenges to (P3) come to mind. The first is that the kind of dependence between the application of concepts and the satisfaction of their application conditions is *not* the kind that determines what's fundamental. The second is an example of a concept that picks out something just as fundamental as the features named in its application conditions. I'll consider each in turn.

Is the kind of dependence I've described the kind that determines what's fundamental? Does being bald depend on the number of hairs one has in the same way that the existence of a table depends on the existence of some wood? Contrast the following two dependence relations. First, the concept of **being a table** *applies* to something only if it has a hard, flat top. Second, a table *exists* only if some hard, flat-

topped wood exists. At least the second of these – the one that makes claims about the existence of objects – is definitely related to questions about fundamentality. Both of these say something about tables and their tops, and both describe a kind of dependence. However, the first describes a dependence between a concept's application and the satisfaction of its application conditions. The second describes a dependence between the existence of an object and the existence of its materials.

What connects these two kinds of claims to fundamentality is that the relationship between the existence of a table and the application of the concept of **being a table** to an object is intimate. Tables exist if and only if there are things to which the concept of **being a table** applies. Bald people exist if and only if there are people to whom the concept of **being bald** applies. Since facts about the existence of objects determine which things are fundamental, but objects exist if and only if the relevant concepts truly apply to them, then facts about the application of concepts also determine which things are fundamental. Thus, there is good reason to think that the kind of dependence between the application of concepts and the presence of the features named by them is the kind relevant to fundamentality.

The second objection to (P3) says that there is a concept – whose application conditions permit non-trivial description – that picks out something just as fundamental as the features named in its application conditions. The challenge is to say which concept that is. Here's what the concept must be like. There must be some description of the conditions that determine whether the concept applies, and the description can't

just be a metaphysical circuit. So the concept of **existing** doesn't work, because the condition that something be *real* amounts to no more than that it exists.[8] Nonetheless, whether the feature the concept picks out is present does *not* metaphysically depend on whether the features named in the application conditions are present.

Since all concepts metaphysically depend on the satisfaction of their application conditions, I can't see how the needed concept is possible. How can any feature be just as fundamental as another, while at the same time being metaphysically dependent on the other, unless the metaphysical dependence is somehow trivial or circuitous? And if this dependence be trivial or circuitous, then how can it be captured by a non-trivial description of any application conditions? It would be like saying that whether the concept of **being bald** applies to someone depends on how many hairs one has, but that being bald doesn't depend on the number of hairs one has.

(P4) says that if the features that determine whether a concept applies are more fundamental than what the concept picks out, then whatever the concept picks out isn't fundamental. This just follows from what it is to be fundamental. No object is fundamental if something else is more fundamental. You can only be fundamental by being metaphysically independent, dependent on nothing for your existence. But if something is more fundamental than you, then it depends on fewer things than you. You can't depend on a relatively greater variety of things while also being totally independent. Thus, (P4) is true by definition.

---

[8] Remember that nothing important turns on the claim that there is a genuine concept of **existing**. For my purposes, this example is purely illustrative. See footnote 5.

This exhausts the possible ways of rejecting each premise of my argument. (P1) is assumed, (P4) is true by definition, there is no sound counter example to (P3), and rejecting (P2) requires serious metaphysical commitments. However, sometimes a conclusion is so implausible that it seems one of the premises must be false, despite the fact that it's difficult to say which. I will now consider possible reasons for thinking this is the case with regard to the Main Argument.

§2.1 – Counter-Intuitive Implications and Theoretic Neutrality

Are there any implications of my argument's conclusion that are themselves so counter-intuitive, they convince us that *something* must be wrong with my argument no matter how hard it is to pinpoint? One species of such an implication would be an apparent counter example to the conclusion. Counter examples to my arguments' conclusion are cases where a concept is vague but nonetheless seems to pick out a fundamental feature, or where a concept picks out the fundamental but nonetheless seems vague. Of course, there might be all kinds of concepts that at first seem to be vague and also pick out fundamental features, but once one appreciates the intuitive strength of my argument, we feel it's better to revise our initial intuitions. §5 and §6 will be entirely about such cases. Here I am only concerned with the possibility of decidedly counter-intuitive implications – or cases in which I do not think our intuitions are open to revision, even if my arguments are very strong. One candidate is the concept of **being fundamental** itself. Does the concept of **being fundamental** pick out a fundamental feature? Does the application of the concept of **being fundamental** depend on nothing further, like the application of the concept of **existing**, or does it apply by virtue of

[17]

some more fundamental features? If the former, then my argument entails that the concept of **being fundamental** cannot be vague. The worry is that the concept of **being fundamental** *is* vague.

I will consider this possibility in depth, mostly because the concept of **being fundamental** does show some of the symptoms of vagueness. For example, fundamentality comes in degrees. Some things are more fundamental than others, and one can list things in order approaching absolute fundamentality. Tables are not very fundamental, arrangements of wood are a little more fundamental, arrangements of molecules are a little more fundamental, and so on. Of course, this is not sufficient to show that the concept of **being fundamental** is vague. **Fundamentality** is vague in the sense relevant to my argument only if its vagueness can be demonstrated using a sorites series. Its vagueness can be demonstrated in a sorites series only if one can provide a non-trivial description of its application conditions to use when describing the various cases.

Here is one attempt at describing non-trivial application conditions for the concept of fundamentality: something is fundamental only if it is metaphysically independent. In other words, something is fundamental only if it does not depend on anything else whatsoever. This description does not explicitly use the concept of fundamentality, so it is not obviously trivial. It also permits us to describe cases that have the feel of a sorites series. For example, in a_1 whether the feature obtains depends on nothing at all. This feature is clearly fundamental. In a_2, whether the feature obtains depends on only one thing. This feature is less fundamental. In a_3, whether the feature obtains depends on two things – and so on. This is structurally

reminiscent of the baldness series, where in a_1 the man has no hairs, in a_2 he has one hair – and so on.

However, I have not claimed that if non-trivial descriptions of application conditions *can* be given for a concept, then one *can* construct a sorites series demonstrating its vagueness. I have only claimed that if one *can't* give such a description, then one *can't* demonstrate the concept's vagueness. So while the concept of **being fundamental** seems to boast a non-trivial description of its application conditions, that doesn't mean it can be featured in a genuine sorites series. Indeed, in the example above, we can see just why it is not a genuine sorites series, even though it looks like one on the surface. The only case that can possibly satisfy the metaphysical independence condition is a_1. Even though the feature in a_2 is *more* fundamental than the feature in a_3, it is still clearly not fundamental.[9] Thus, just because a concept's application conditions can approach satisfaction in increments over a series of cases doesn't mean that the concept will produce a sorites series.

---

[9] The point that sorites series using the concept of **being fundamental** seems impossible is echoed by Michael Perkins: "…there are no borderline or vague cases of fundamentality. Precise Fundamentality is extremely plausible. At any rate, the usual test for vagueness strongly suggests that Precise Fundamentality is true. It appears to be impossible to construct a 'sorites series' for fundamentality—a series of properties <F1…Fn> such that:
  - Each Fi differs ever-so-slightly from Fi+1;
  - F1 is determinately fundamental;
  - Fn is determinately not fundamental;
  - There is no 'determinate cut-off point' in the series—i.e., no property Fj such that Fj is determinately fundamental and Fj+1 is determinately not fundamental.
I simply cannot envisage any such sorites series. For this reason, there is a strong presumption that Precise Fundamentality is true." (2012, 98-99).

This is also suggestive that the description of **being fundamental** that I've given isn't really non-trivial. Remember that it's not enough for the description of the application conditions to avoid employing the concept *explicitly*. The description of the application conditions must avoid *any* circularity that results from using the concept. For example, I gave the following toy description of the application conditions for the concept of **existing**: something exists only if it is real. While this description does not employ the concept of **existing** explicitly, it is still trivial, since ultimately something is real only if it exists (or something to that effect). The concepts employed in my toy description don't explicitly include the concept being described (the concept of **existing**), nonetheless they *employ* the concept of **existing**. Is there any good way to understand this description unless being real *just is* existing? Consider the suggested condition for fundamentality: some object/feature is fundamental only if it is metaphysically independent. But under what conditions is something metaphysically independent? What determines the extent to which an object depends for its existence on others?

Answering this question is probably just as hard as giving a non-trivial description of the application conditions for the concept of **existing** – which is to say very hard. Considering two possible answers will be enough for my purposes. First, perhaps *nothing* determines whether an object depends for its existence on others. On this suggestion, metaphysical independence is a fundamental feature. However, it seems strange that being metaphysically independent would be more fundamental than being fundamental. This strangeness suggests a second answer. The reason that metaphysical independence and fundamentality seem like equally fundamental features is that they are identical. There is really only one feature, and it goes by

[20]

different names. Contrast this with concepts whose application conditions permit genuine, non-trivial descriptions, like the concept of **being bald**. It's not that being bald is the same as having no hair. It's just that having no hair determines that one is bald. In that case, "having a certain amount of one's scalp showing" is not just another name for being bald. These are distinct features, one of which determines the presence of the other. Being fundamental, on the other hand, *just is* being metaphysically independent. I'll call this view the identity picture: sometimes, it sounds like we can sufficiently describe concepts in terms of other concepts, but really all we do is give *one* concept several names. The identity picture permits one to say that metaphysical independence is *not* a genuine condition for fundamentality. Metaphysical independence just is fundamentality. Thus, the metaphysical independence condition does not allow us to generate a sorites series for the concept of **being fundamental**.

This issue is deep, and probably deserves a dissertation of its own, so I don't intend to decisively settle it here. However, I can give one good reason for accepting the identity picture that motivates my use of it. One independent reason to accept the identity picture is that it has broad explanatory power. The concept of existence should pick out something just as fundamental as the concept of being real. The feature of being real can't be any more fundamental than the feature of existence. The identity picture explains why that would be, and produces the correct results. Existing *just is* being real. There is only one concept, and we give it several names. There is only one feature, and whether it obtains does not depend on anything. This makes sense of many of our struggles in conceptual analysis, where we feel we can *rename* concepts, and yet we cannot sufficiently *describe* them.

The result of this discussion is our certainty that the concept of **being fundamental** is not vague. I gave the only suggestion I could imagine for producing a sorites series using the concept of **being fundamental** – but it fails for two reasons. First, just because one can give a non-trivial description of a concept's application conditions doesn't mean that one can demonstrate its vagueness using a sorites series. Second, an independently plausible view entails that my proposed description of the application conditions of **being fundamental** is trivial after all. My thesis only entails anything about fundamentality if the concept of **being fundamental** is vague. Unless one can demonstrate that the concept of **being fundamental** is vague using a sorites series, the concept of **being fundamental** is not a genuine counter example to my thesis.

So far I've argued that one way we can tell whether a concept is vague is by investigating whether the concept picks out something fundamental. If so, then the concept is not vague. Does my argument commit me in any way to a particular theory of vagueness, although I have tried to assume only general features of vagueness? Is it unacceptable to proponents of one theory or another?

The two standard views about the nature of vagueness are vagueness-as-indeterminacy and epistemicism. The indeterminist's thesis says that vagueness is a kind of indeterminacy. The epistemicist's thesis says that vagueness results from in-principle-unknowable facts. There are other serious views about the nature of vagueness, but I will only consider these two.

First, consider the theory that vagueness is a kind of indeterminacy. This view comes in two varieties. The first is what I will call the conceptual indeterminacy view. The conceptual

indeterminacy view makes a generalization about the application conditions of vague concepts. It says that they clearly tell us that the concept applies sometimes, that the concept does not apply other times, and no more than that. For example, it might be that the concept of **being bald** applies to an object when it has any number of hairs less than 2,500, and that it does not apply to an object when it has any number of hairs greater than 3,500. However, nothing in the concept's application conditions tells us whether the concept applies to objects with between 2,500 and 3,500 hairs. On the conceptual indeterminacy view, it is vague whether a man with 3,000 hairs is bald because the application conditions for the concept of **being bald** themselves do not determine whether it applies to a man with 3,000 hairs.

The conceptual indeterminacy view may seem inconsistent with the claim that if a concept picks out a fundamental feature, then it is not vague. On the conceptual indeterminacy view, it might be that all of our vague concepts are ultimately built using only clear concepts and clear logic. It's just that they have gaps in their application conditions that render us powerless to apply them in some cases. It's possible to analyze the concept of **being bald** in terms of numbers of bosons and their space-time locations. The reason the concept of **being bald** is vague is not that it's built using some vague concepts or vague logic. The reason is that its application conditions are sometimes silent about whether the concept applies. As a result, the fact that a concept picks out a fundamental feature does not entail that the concept is clear. According to the conceptual indeterminacy view, a concept might be such that it indirectly picks out fundamental features, but has gaps in its application conditions such that the concept is vague.

Here are three responses. First, if conceptual indeterminacy is *any* kind of vagueness, it is just not the kind that I have in mind. The kind that I have in mind is explicitly the kind that gives rise to sorites series, and conceptual indeterminacy does not. For example, it is hard to see how the concept of **being bald** can give rise to a sorites series when its application conditions are indeterminate this way. For every case in the series a_1-a_{2500}, the concept clearly applies. For every case in the series a_{3500} and on, the concept clearly does not apply. Beginning clearly with case a_{2501}, it is impossible to say anything about whether the concept applies. The concept's application conditions simply do not say. It is very clear, given this picture of the concept of **being bald**, that we should not say that the concept applies in case a_{2501}. There is a precise boundary to the application of the concept, which is exactly *not* how a sorites series works. A sorites series ratchets up the fuzziness of the concept's application in each case, so that there is no precise boundary *at all*. Thus, this view of vagueness is simply not within the scope of my project. It's not that I have assumed that the view is false for the sake of my argument, so much as that it concerns something that my argument does not.

Second, even if it were a view about the nature of vagueness of the kind that I have in mind, there is independent reason to think that it's false. It does not capture what we intuitively mean when we say that a concept is vague. When we consider the concept of **being bald**, we typically are not sure whether to say the middle men in the series are bald or not bald. We somehow get the sense that they are sort of…both, yet sort of neither. When we are given the above analysis of the concept of **being bald** though, it is obvious that we should not say anything of the middle men. We do not feel uncomfortable saying the middle men are bald.

[24]

Rather, it is plain that if we are asked whether the middle men are bald, the correct response is silence. Surely this is not the phenomena we describe when we refer to vagueness.[10]

Third, this theory of vagueness is necessarily false, since it says that some concepts lack necessary application conditions. It says that vague concepts are such there are some conditions sufficient for the concept to apply and some conditions sufficient for the concept not to apply. But what's needed to be bald? Well, if one has fewer than 2,501 hairs, that's enough to be bald. But it's not that one *must* have fewer than 2,501 hairs, since the only condition under which the concept does *not* apply is when one has 3,500 hairs. How can this be? Does it even make sense? Surely any genuine concept is such that there are at least some necessary conditions for its application.

Lastly, since the only things I have assumed in my argument are general features of vagueness that any theory would endorse, we may just take disagreement as evidence that conceptual indeterminacy is an inadequate view. If that view does not square with totally general features of vagueness, independently of any theory, then so much the worse for it. Thus, the conceptual indeterminacy view of vagueness is not problematic for my argument. First, it does not even seem relevant to my argument, since it isn't the kind of vagueness important to a sorites series. Second, even if it were relevant, it is intuitively false. Third, it is worse than counter-intuitive; it is necessarily false given a plausible view about the general nature of concepts. Lastly, even if it were not counter-intuitive, my argument only employs

---

[10] We might also take this point to be more evidence that the conceptual indeterminacy view is not really concerned with sorites style vagueness – the kind of vagueness that my project is concerned with. For a sustained argument for this conclusion, see Barnett (2000) and (2009), and Sainsbury (1991) and (1996).

general features of our notion of vagueness. If the view is in tension with even general features of vagueness, it is not very likely to be a good view anyway.

The second view that says vagueness is a kind of indeterminacy is the Lewisian view. However, on the Lewisian view, vagueness appears at the level of language, rather than at the level of concepts. A word is vague whenever language users haven't settled which concept is the meaning of the word, that is, which set of application conditions the word employs. For example, it's not settled whether a man with 25% of his head sparsely covered with hair is bald, because people who use the word "bald" haven't decided whether 25% coverage counts. We haven't settled – or determined – the rules specifically enough to say whether having 25% of one's head covered in sparse hair is appropriately described by the word "bald."

Since the Lewisian view concerns words rather than concepts, it's not obvious that it competes with my argument any more than the conceptual indeterminacy view. My conclusion can be restated so that it is compatible with the spirit of this view, but it requires terminology I introduce in the next section. Basically, proponents of this view think that *all* concepts pick out fundamental features one way or another. Fundamental features are really the only features that concepts can pick out; they are the only features there are. Thus, my conclusions are compatible with the Lewisian view. Counter examples to my main thesis on that view are impossible, since the consequent can never be false. If a concept picks out a fundamental feature, then it can't be vague, since *concepts* are never vague. I will revisit this point in §4.

There is no conflict between the Main Argument and the epistemicist's view. Consider the baldness series. The epistemicist thinks that there really is a line in the series between the

last bald man and the first man that isn't bald.   For example, perhaps case a_{3499} is the last case

in which the man is bald, and a_{3500} is the first case in which he is not bald. However, there's

something about the concept, or our language, or the world, which makes it *impossible* to know

this about baldness. My argument says nothing about how we know whether the concept of

baldness applies to any case in the series. I've claimed that the application of vague concepts is

determined by the satisfaction of their application conditions. This is neutral with respect to

whether there is always an answer about when they apply, and it is neutral with respect to

whether we can know the answer.[11]

§3 – The Reverse Thesis

Thus far, I have defended the thesis that if a concept is vague, then what it picks out is

not fundamental. However, I have not defended the reverse claim that if a concept doesn't

pick out something fundamental, then it is vague. Together, these claims form the stronger

claim that fundamental features are picked out by all-and-only clear concepts. The stronger

claim is attractive since it would permit us to translate back and forth between our intuitions

about which concepts are vague and our intuitions about which objects and features are

fundamental. In this section, I will try and establish the reverse claim.

Many concepts pick out non-fundamental things. In contrast, hardly *any* concepts pick

out fundamental things. Clear and vague concepts are also asymmetrical in this way. Many

---

[11] For an extensive defense of epistemicist views of vagueness, see Williamson (1994).

concepts are vague; hardly any concepts are clear.[12] This provides two kinds of evidence.[13]

First, since most concepts are vague and most concepts fail to pick out the fundamental, it's no

surprise that there's lots of overlap between them. For example, the concepts of **being bald**,

being tall, and **being childish** all pick out features that aren't fundamental. All of these

concepts are *also* vague. On the other hand, many clear concepts overlap with concepts of

fundamental things. The rare concepts that satisfy either category often satisfy both. For

example, the concept of **existing**, of **being possible**, and of **being a number** are all clear. All of

these concepts *also* pick out fundamental features of reality. Thus, we get evidence for the

reverse thesis of both kinds: overlap between concepts of non-fundamental things and

concepts that are vague, and overlap between concepts that are not vague and concepts of the

fundamental.

Nonetheless, the reverse thesis is false as stated. A family of counter examples

demonstrates this, but a distinction is needed to generate them. One can group concepts

according to their analyzability. Unanalyzable concepts are to be contrasted with analyzable

concepts. We cannot give non-trivial descriptions of application conditions for unanalyzable

concepts. For example, the concepts of **being fundamental** and of **being a number** are

unanalyzable. In contrast, the concept of **being fundamental and a number** is analyzable. The

---

[12] This may not be the right way of describing the asymmetry if there are actually an infinite
number of both concepts. Since it doesn't make sense to divide infinite quantities using ratios,
the asymmetry would instead appear among the concepts we *use*, or the concepts we are
aware of, or something to that effect.
[13] I recognize that this kind of double-evidence may result in the raven's paradox, but I have no
space to consider that here. For my purposes, I'm happy to use evidence of both kinds. I don't
think that the clear/fundamental is as disparate from the vague/non-fundamental as the
black/raven is from the non-black/non-raven.

concept of **being fundamental and a number** is an analyzable concept that can be analyzed in terms of two unanalyzable concepts (**being fundamental** and **being a number**). When a concept is built out of other concepts this way, I'll call the concepts it's built out of the constituent concepts. So in the above case, the concepts of **being fundamental** and **being a number** are the constituent concepts of the concept **being fundamental and being a number**.

The concept of **being fundamental and being a number** picks out the fundamental in one sense. Its constituent concepts all directly pick out a fundamental feature. However, in another sense the concept does not pick out the fundamental. Whether the concept of **being fundamental and a number** applies to something *depends* on two things: whether the object is fundamental, and whether the object is a number. If something is fundamental, then it does not depend for its presence on anything else. But analyzable concepts apply to an object only if their constituent concepts apply to the object. Whether an analyzable concept applies to an object *depends* on whether its constituent concepts apply to the object. Thus, in the important sense, analyzable concepts do not pick out fundamental features or objects.

Despite the fact that analyzable concepts do not pick out fundamental features, they can still be clear. We could not use a sorites series to demonstrate that the concept of **being fundamental and being a number** is vague, because it would be clear for every case in the series whether the concept applied. Supposing that being a number is a fundamental way of being, the concept of **being fundamental and being a number** clearly applies to every number and nothing else. Thus, although the concept of **being fundamental and being a number** is not vague, it does not directly pick out a fundamental feature.

[29]

Indeed, we could generate an infinite list of counter examples to the reverse thesis this way. By applying clear logical operators to the list of clear concepts, we could indefinitely generate a list of clear concepts that do not directly pick out fundamental features. For example, the concept of **being fundamental or being a number**, of **existing and being a number**, of **existing and being fundamental and being a number**, and so on.

The problem is that we need a principle that bridges clearness and fundamentality. Indeed, there are many notions that seem *somehow* related to clearness and fundamentality: simplicity, irreducibility, bruteness, and so on. I've chosen "unanalyzable" arbitrarily from this list. Thus, the principle I will now consider says that if a concept is clear, then it's unanalyzable. The concept of **being fundamental and being a number** serves as a counter example to this principle. But the same kinds of counter example will undermine the principle no matter which related notion it employs. Analyzable concepts are reducible to their constituents, they are complex rather than simple, and so on. Since the concept of **being fundamental and being a number** is clear, clearness does not entail anything about fundamentality or the other related notions.

That said, I think that these counter examples are highly artificial; they speak to the letter of the reverse thesis, but not its spirit.[14] As I suggested at the start of this section, we have *some* reason to expect a deep connection between clearness and fundamentality. There also seem to be several intimately related notions that might connect the two. The analyzable-but-clear concepts show that clear concepts need not be fundamental in the relevant sense.

---

[14] Thanks to my colleague Amanda Brockman for emphasizing this point in an earlier draft.

However, they do so only by clearly combining clear concepts that *are* fundamental in the relevant sense.

This suggests that the reverse thesis can be revised to avoid these problems while retaining its spirit. If a concept is not vague, then either (i) what it directly picks out is fundamental, or (ii) the concept is ultimately a product of applying clear logical operators to concepts that directly pick out fundamental features. For example, the concept of **being fundamental and being a number** is analyzable. It does not directly pick out anything fundamental. However, it's the product of conjoining concepts that do: the concept of **being fundamental** and the concept of **being a number**. It's important that the revised principle says "ultimately," because some analyzable concepts have other analyzable concepts as constituents. For example, the concept of **(being fundamental and being a number) and (existing and being necessary)** has the analyzable concept of **being fundamental and being a number** and the concept of **existing and being necessary** as constituents. But each of those constituents is analyzable. For example, the former has the concept of **being fundamental** and the concept of **being a number** as its constituents. Thus, the concept of **(being fundamental and being a number) and (existing and being necessary)** is ultimately a product of concepts that pick out fundamental features. The concepts at the bottom of this conceptual pyramid pick out fundamental features. Henceforth, I'll call these bottom-level concepts "ultimate concepts."

The revised reverse thesis says that if a concept is not vague, then it either picks out a fundamental feature directly (like the concept of **being fundamental**) or its ultimate

[31]

constituents pick out fundamental features directly (as with the concept of **being fundamental and a number**). This formulation avoids the counter examples that refute the reverse thesis.

The clear-but-analyzable concepts that plagued the original reverse thesis all satisfy condition (ii) of the revised reverse thesis. They do not directly pick out fundamental features, but they are ultimately built up out of concepts that do.

Even if the revised reverse thesis escapes one serious problem, are there compelling reasons to endorse it? To begin with, I think that this turns out to be a fairly common intuition. In §5 and §6, I will try to provide readings of a variety of philosophers that make this intuition plain. For now I'll set this aside and attempt to provide an independent argument for the revised reverse thesis. As I have described ultimate concepts, they do not admit any non-trivial analysis. They are unanalyzable, and we cannot describe their application conditions without appealing to the ultimate concepts themselves. For example, the concept of **being fundamental** is an ultimate concept. It is the sort of concept at the bottom of our conceptual pyramid that cannot be usefully analyzed in terms of other more basic concepts. We can rename it, but we cannot describe non-trivial conditions under which it applies.

This claim, together with two important features of vague concepts entail that there are only two kinds of clear concepts: the ones that are themselves ultimate concepts, and the ones that are ultimately built out of those concepts using clear logical operators. The first important feature of vague concepts is that they are all analyzable. As I argued in §1, we cannot demonstrate a sorites series for any unanalyzable concept. If no ultimate concepts are

analyzable, but all vague concepts are analyzable, then no ultimate concepts are vague. In other words, since all ultimate concepts are unanalyzable, it is always clear when they apply.

The second important feature of vague concepts is that their vagueness is infectious. What I mean by this is that there's *some* connection between the vagueness of a concept's constituents and the vagueness of the concept itself. Bertrand Russell at least thought so, which partly led to his position that *all* language is vague:

> And the vagueness of the knowledge derived from the senses infects all words in the definition of which there is a sensible element. This includes all words which contain geographical and chronological constituents, such as "Julius Caesar," "the twentieth century," or the "solar system." (1923, 87).

But other philosophers have had the same intuition to less radical effect. For example, Michael Antony can be read as alluding to it in his (2008). He gives the following argument that the concepts employed by functional theories of mind must be vague:

> One is to focus on concepts for the inputs and outputs in the functional definition of F…the concepts for most if not all such inputs and outputs will be vague. Such vagueness, however, will infect F, since F is defined in terms of such concepts. (2008, 527).

The relevant point here is this: functional concepts have particular input and output concepts as constituents. Functional concepts must be vague *because* the concepts that pick out particular input and output types are vague. The vagueness of the constituents infects the

vagueness of the concept itself. The intuition is also at work in Achille Varzi's (2003). She points out that many vague concepts inherit their vagueness from vague constituents: "…certainly 'married to a bald man' is truly a vague predicate, even though its vagueness is entirely parasitic upon that of 'bald'" (2003, 296).

We should *not* characterize this intuition as follows: for any concept, if one of its constituents is vague, then the concept is vague to at least that extent. First, the notion of being vague to an extent is opaque. For example, we might naturally cash this out in terms of the number of borderline cases that a concept admits. If, for some concept, there is only one possible case in which it is vague whether the concept applies, then that concept is less vague than if there were only two possible cases in which it was vague whether the concept applied. The problem is that our ordinary vague concepts admit countless borderline cases. For example, the concept of **being bald** has as many borderline cases as there are possible arrangements of fundamental particles in the range between roughly 25% hair coverage and roughly 75% hair coverage. The concept of **being a hair** has as many borderline cases as there are possible arrangements of fundamental particles in the range between clearly hair-structured and clearly not-hair-structured. Thus, the concepts of **being bald** and **being a hair** are vague to the same extent. But this seems wrong. We naturally think of **being a hair** as less vague than **being bald**. Presumably, this is what inspires people to say that it is *less* vague whether one has a hundred hairs than it is whether one is bald.

Furthermore, I don't think that counting borderline cases gets at the heart of the intuition anyway. The intuition is that a parent concept will have some borderline cases that

[34]

are *identical* to borderline cases of its constituents. A parent concept inherits its vagueness from its constituents in the following sense: some of the borderline cases of its constituents translate to borderline cases of the parent concept. For example, consider Antony's case. Suppose that functional state F is defined by taking input I and producing output O. Suppose that in a sorites series, o_{5000} is clearly a borderline case of O. Then in a distinct sorites series, there will be some case f_x, such that it is vague whether F applies in f_x *because* the status of O in that case matches o_{5000}. It's vague whether I am conscious in case f_x because it is vague whether my brain is producing sufficient electrical current. The way in which it is vague whether my brain is producing sufficient electrical current corresponds to the description in o_{5000}.

Or, consider Varzi's case. Suppose that Mary is married to Harry, the archetype of vagueness with respect to **being bald**. It is vague whether the concept of **being married to a bald man** applies to Mary's situation *because* it has the concept of **being bald** as one of its constituents. Harry is clearly a borderline case of **being bald**. Thus, Mary's situation is such that it is vague whether **being married to a bald man** applies, *because* it is vague whether **being bald** applies to the man she married**.**

Still, cashing the intuition in terms of matching cases isn't quite right either. The concept of **being married to a bald man** doesn't have a perfectly matched borderline region with the concept of **being bald**. Mary could have instead fallen for David, who is clearly bald. Yet Mary's situation could still be vague if we catch them at the right moment of their vows. David has begun to say "I do," but it's vague whether he's said it yet. This situation is a

borderline case of **being married to a bald person**, but *not* because it is a borderline case of

being bald. David is clearly bald.

A conservative statement of the intuition might sound more like this: for any concept, if

it has any vague constituent, the concept will have at least one borderline case that is also a

borderline case of its constituent. Even this is not correct though. As Lewis points out,

"sentences with vague constituents are not necessarily vague: 'This is cool or warm, but not

both' is true at all delineations…" (1970, 64). Indeed, Roy Sorensen basically provides a formula

for building clear concepts out of vague constituents and clear operators in his (1988).

> The first class…relies on the precision of one of the defining clauses. Since 'self-
>
> identical' applies to everything, it is perfectly precise. Disjoining it with any vague
>
> expression leaves its precision intact. Just as everything is in the extension of
>
> 'self-identical', every- thing is in the extension of 'is self-identical or a heap of
>
> sand'…Conjoining any vague term with a contradictory expression such as
>
> 'round-square' or 'largest prime' will not expand their extensions. Just as nothing
>
> falls under 'round-square', nothing falls under 'large round-square'. (1988, 270)

If a concept is such that it clearly applies to everything, or it clearly does not apply to anything, I

will call it a *perfectly* clear concept. Thus, clear disjunctional concepts can have vague

constituents (as long as at least one disjunct is perfectly clear), and clear conjunctive concepts

can have vague constituents (as long as one conjunct is perfectly clear). So how should we

characterize the intuition that vagueness is infectious?

We might characterize it by exempting Sorensen style counter examples. For example, for any concept, if one of its constituents is vague, and none of its constituents are perfectly clear, then the concept admits at least one borderline case that is a borderline case of its vague constituent.

This entails that no clear concept has any vague constituents, since clear concepts are concepts that are not vague *at all.* If a concept has even one vague constituent at any level of its construction, then the concept has *at least* one borderline case. In particular, it has a borderline case that overlaps a borderline case of its vague constituent. But even one borderline case precludes clearness. So no clear concept has any vague constituents. Thus, all clear concepts that are not ultimate concepts have only clear concepts as constituents. If it were otherwise, then clear concepts could have borderline cases, which is impossible. But the only concepts that do not have ultimate concepts as constituents at some level of their construction are ultimate concepts themselves.

Together with the fact that no ultimate concepts are vague, this entails that there are only two kinds of clear concepts: the ones that are themselves ultimate concepts, and the ones that are ultimately built out of those concepts using clear logical operators. For example, the concept of **existing** and the concept of **being fundamental** are each clear concepts that are ultimate themselves. The concept of **existing and being fundamental** is a clear concept that is ultimately built out of those concepts using clear logical operators. Since no ultimate concepts admit analysis, all clear concepts are either (i) unanalyzable or (ii) have unanalyzable concepts as their ultimate constituents.

Consider the concept of **being fundamental**. It is an ultimate concept. It is the sort of concept at the bottom of our conceptual pyramid that cannot be usefully analyzed in terms of other more basic concepts. I argued in §2 that the only obvious analysis of the concept of **being fundamental** is actually just another name for the same concept: the concept of **being metaphysically independent**. But why is this? What is it that makes some concepts unanalyzable?

When we give substantive, non-trivial analyses of concepts, we do so by giving non-trivial descriptions of the concept's application conditions. These kinds of description capture a dependence between the features or objects the concepts pick out and the features or objects that we refer to in their application conditions. Since substantive analyses are descriptions of the dependence between what concepts pick out and what their application conditions refer to, being unable to give an analysis means there is no dependence to describe. But if there is no dependence between what a concept picks out and whatever we might refer to in its application conditions, then the concept picks out something fundamental. So, unanalyzable concepts pick out fundamental features or objects.

Thus, the revised reverse thesis follows from the nature of ultimate concepts, the way that vague concepts work, and how the possibility of analysis relates to fundamentality. All clear concepts either directly pick out the fundamental, or are ultimately built out of concepts that directly pick out the fundamental. Here is the argument in an explicit premise by premise format.

The Reverse Argument

[38]

P1. If a concept is not vague, then it is either (i) ultimate itself, or (ii) ultimately built using only clear concepts and clear operators.

P2. All ultimate concepts are unanalyzable.

P3. So, if a concept is not vague, then it is either (i) unanalyzable, or (ii) ultimately built using only unanalyzable concepts and clear operators.

P4. If a concept is unanalyzable, then what it picks out is fundamental.

P5. So, if a concept is not vague, then either (i) what it picks out is fundamental, or (ii) it is ultimately built using only clear logical operators and concepts that pick out the fundamental.

§4 – Potential Problems

(P2) is the definition of an ultimate concept. (P3) follows from (P1) and (P2) by substitution. (P1) and (P4) are the only premises to defend. In this section I will consider objections to each. I argued above that (P1) follows from the nature of vague concepts. No clear concept has *any* vague constituent. As I have stated throughout this paper, I do not mean to assume any particular theory about the nature of vagueness. Indeed, at the end of §2 I tried to show that my Main Argument is perfectly neutral about the particular nature of vagueness. Thus, when I say that (P1) follows from the nature of vagueness, I mean to say that it follows from general properties of vagueness that any view on the nature of vagueness would acknowledge. However, (P1) may seem in tension with the Lewisian view of vagueness I described at the end of §2.1.

The Lewisian view is that there aren't any vague *concepts*. Rather, there are vague

words. Words are vague when it is indeterminate which of a set of clear concepts we appeal to

when we use the word. For example, our word "bald" is vague, because our linguistic

community has not settled which of the following clear concepts that word employs: the

concept of **having fewer than 3,000 hairs**, the concept of **having fewer than 3,001 hairs**, the

concept of **having fewer than 3,002 hairs**, and so on.[15] It is indeterminate which of those we

mean by "bald", but for every possible answer, the relevant concepts are clear.

One may think that (P1) is incompatible with this view, because on the Lewisian view

(P1) does not really distinguish between clear and vague concepts. The Lewisian about

vagueness might concede that all clear concepts are either ultimate themselves, or are

ultimately built out of clear concepts and clear logical operators. Yet the Lewisian about

vagueness would say that the concepts I'm calling vague concepts are like that too. The

concept of **being bald** according to the Lewisian is ultimately constructed using only clear

constituents and clear logical operators. It's just that it's indeterminate which set of clear

constituents and logical operators we mean to employ when we use the *word* "bald," and thus

vague which concept is the concept of **being bald**. Vagueness appears at the level of language,

not of concepts. If it is vague whether someone is bald, it is for this reason: although there are

several candidate concepts, all of which either clearly apply or clearly do not apply to the

person in question, it is indeterminate which one we mean by our word "bald."

---

[15] In §1 I said that the concept of **having exactly *n* hairs** is vague too, since the concept of **being a hair** is vague. In this section I am just speaking shorthand. One should really substitute the concepts I suggest are being employed by our word "bald" for concepts that have exclusively ultimate constituents.

On the Lewisian view, it is true that (P1) does not really distinguish between the concepts that I have called vague – like the concept of **being bald** – and the concepts I have called clear – like the concept of **existence**. However, (P1) is technically not false on the Lewisian view. Technically, it is *necessarily* true, since on that view there are no vague concepts. In order for (P1) to be false, some clear concept would have to be such that it was neither ultimate, nor ultimately constructed out of clear concepts and logical operators. But *all* concepts are either ultimate themselves or have such concepts as their ultimate constituents. It's just that sometimes it's indeterminate which clear concept we mean to use. Another way of putting this is that on the Lewisian view, (P1)'s contrapositive has a necessarily false antecedent. If a concept is not either (i) ultimate itself, or (ii) ultimately built using only clear concepts and clear operators, then it is vague. But there are no such concepts.

Thus, (P1) is compatible with the Lewisian view. However, the Lewisian view would not characterize our ordinary vague concepts the way that I have. A Lewisian proponent of my argument might instead say that if a word is not vague, then either (i) we determinately use it to employ an ultimate concept, or (ii) we determinately use it to employ a concept that is ultimately built using only clear concepts and clear operators. Indeed, the entire argument can be rephrased in terms of words that are vague and the concepts that we mean by them.

The Reverse Argument for Lewisians

P1. If a word is not vague, then either (i) it determinately has an ultimate concept as its meaning, or (ii) it determinately has a concept that is ultimately built using only clear concepts and clear operators as its meaning.

P2. All ultimate concepts are unanalyzable.

[41]

P3. So, if a word is not vague, then either (i) it determinately has an unanalyzable concept as its meaning, or (ii) it determinately has a concept that is ultimately built using only unanalyzable concepts and clear operators as its meaning.

P4. If a concept is unanalyzable, then what it picks out is fundamental.

P5. So, if a word is not vague, then either (i) it's determinate that the word's meaning is a concept that picks out something fundamental, or (ii) it's determinate that the word's meaning is a concept that is ultimately built using only clear logical operators and concepts that pick out something fundamental.

After making Lewisian revisions to the Reverse Argument, its conclusion almost seems trivial. If a word is clear, then it determinately employs one concept – either an ultimate concept or a concept built only using ultimate concepts and clear logic. But on the Lewisian view, that's just what it is for a word to be clear. There's nothing more exciting about my conclusion than the Lewisian view of vagueness itself.

To see this, consider a debate about whether particular words are vague. Are those words vague, or are they clear? Well, they are vague if there are several candidate concepts we might mean by them, but it is indeterminate which one we really mean. They are not vague if it is determinate which concept we employ with our words. But the only *possible* concepts, according to the Lewisian, are either ultimate or built using only ultimate concepts and clear logic. That's what *all* concepts are like. Thus, if a word is not vague, then we determinately use it to employ either an ultimate concept, or a concept built using ultimate concepts and clear logic, because those are the only concepts *any* clear word could possibly employ. In some

sense then, *all* of our concepts pick out the fundamental, directly or indirectly. It's just that sometimes we haven't settled which set of concepts we mean, and thus which fundamental features we pick out.

I think that this sort of view is attractive to a certain kind of philosopher. This kind of philosopher likes a picture of the world on which all there *really is* are the fundamental objects and features. The sorts of features that we take to be not fundamental are just an illusion – a convenient short hand, or useful deception. We just imagine them overlaid onto the *real* world because it is useful, or because we are lazy, or because we are confused, or because it is fun. I will make no more comment about the virtues or vices of this view. For my purposes, it is enough that my argument is not really in tension with the Lewisian view. In fact, the argument can be rephrased in such a way that the Lewisian would happily accept it as a consequence of their theory of vagueness.

Objections to (P4) involve a very difficult discussion that will take me far afield of my actual topic, so I will treat the whole category briefly. (P4) says that if a concept is unanalyzable, then what it picks out is fundamental. In other words, the reason we can't provide a substantive analysis of some concepts is that the features they pick out don't depend on any others. You can't describe non-trivial application conditions for the concept because application conditions capture a kind of dependence between features. But when the feature in question is metaphysically independent, there's no dependence to describe.

There are a cluster of views that reject this picture. G.E. Moore argued that our concept of **goodness** was unanalyzable, but that goodness nonetheless supervenes somehow on human

psychological states.[16] Which things our concept of **goodness** applies to depends on which

things make us happy. Yet we cannot analyze our concept of **goodness** in terms of happiness,

or any other terms. David Barnett has argued that our concept of **vagueness** is unanalyzable,

yet vagueness supervenes on our use of the "roughly" operator.[17] Several philosophers have

argued that we cannot analyze our concept of **causation**, yet causation supervenes on facts

about possible worlds.[18]

In each of these cases, we get an argument that any possible analysis of the relevant

concept must fail. For example, Moore's open question argument is intended to show in

principle that absolutely *no* analysis of our concept of **goodness** can be correct.   Barnett just

argues individually against every possible analysis of our concept of **vagueness**. Yet these

philosophers don't think that the features in question are fundamental. We could influence the

application of these concepts by changing the world in other ways. For example, the

application of the concept of **goodness** would be different than it actually is, if the sorts of

things that pleased people were different.

If any such view is correct, then my current statement of (P4) is false. Some concepts

cannot be analyzed, but nonetheless they do not pick out fundamental features. They pick out

features that supervene on other aspects of the world. This introduces difficult issues about

the general possibility of conceptual analysis throughout philosophy. Here is my only word on

the subject. Insofar as we find (P4) very plausible, we have just so much reason to reject the

possibility of unanalyzable concepts that pick out supervenient features. However much reason

---

[16] See Moore (1903).
[17] See Barnett (2009).
[18] See Barnett (2011, 23).

we have to believe (P4), we have that much reason to think that the concept of **goodness** is

analyzable after all. At least, we have that much reason to think that if it is unanalyzable, it is a

genuinely fundamental feature. In short, it's reason to believe that Moore and these others are

just wrong.[19] But suppose that one or all of the imagined views are correct. A weaker, fallback

position than (P4) distinguishes typical concepts from the confusing kinds of concepts that

Moore and Barnett and others argue are unanalyzable. After all, it seems like there is *some*

intimate relationship between a concept's application conditions and the features the concept

picks out. In most cases where an analysis is possible, it's *because* an analysis captures

metaphysical dependence between features, while in cases where one is impossible, it's

because there is no dependence to capture. For example, the concept of **existing** seems

unanalyzable in any terms, and this is *because* whether something exists is a matter of

metaphysical independence. In contrast, the concept of **being bald** seems analyzable in terms

of hair arrangements, and this is *because* whether someone is bald depends on how their hair is

arranged.

Some concepts are unanalyzable because there's no dependency to describe. I'll call

these ordinary concepts. The concepts that represent a problem for (P4) are unanalyzable for

some other reason. I'll call these Moorean concepts. With Moorean concepts, it's not that

there's no dependency to describe. The features that these concepts pick out are not

metaphysically independent in the world; they do depend on some other features for their

presence. Rather, their analysis is impossible in an epistemic sense. There's something about

---

[19] This is by no means serious philosophic evidence, but I think it's still worth noting that none
of these views enjoy popularity in their respective fields today. I conjecture that this is related
to the serious research that's been done on supervenience and fundamentality during the last
century. Perhaps a principle like (P4) is at the back of our collective disciplinary consciousness.

our conceptual framework, or perhaps our limited evidential resources, that makes it so that

we just can't produce any adequate description of the application conditions for Moorean

concepts – and we couldn't no matter what. If it turns out that the Moorean concepts really

are unanalyzable, despite that they do not pick out fundamental features of reality, then my

argument only concerns ordinary concepts.

This is significantly weaker than the thesis I've been defending. However, it is still

interesting, since almost all of our concepts are ordinary in the sense I've described. So it will

be true of almost all of our concepts that if they are unanalyzable, then they pick out a

fundamental feature. Thus, of all our ordinary concepts, the connection between simplicity and

fundamentality will be intimate. I do not think that this is ad hoc, since surely there is

something very strange going on with the Moorean concepts. As I have argued, it is natural to

think that the *reason* we can't describe application conditions for ordinary, unanalyzable

concepts is that the features they pick out seem fundamental. So pre-theoretically it seems

that something is odd about Moorean concepts anyway. However, the weaker position is open

to a serious question: how can we tell when we're dealing with a weird Moorean concept

rather than an ordinary concept that's unanalyzable? I suspect that we can only rely on already

shaky metaphysical intuitions to navigate this territory. In any event, I have tried to motivate

(P4) the best way I know how, but any further discussion would merit a serious philosophic

work of its own.

§5 – Is Morality Vague?

My main question throughout this work has been this: how can we tell whether a concept is vague? I have argued for at least two criteria for deciding whether a concept is vague. The main thesis says that if a concept is vague, then what it picks out is not fundamental. If it does, then the concept is not vague. The revised reverse thesis says that if a concept is not vague, then either (i) what it picks out is fundamental, or (ii) it is ultimately built using only clear logical operators and concepts that pick out the fundamental.   Thus, we can tell whether a concept is vague by examining what the concept picks out. If it picks out something fundamental, the concept is not vague. If it doesn't, then either the concept is vague or its ultimate constituents are clear concepts.

In the rest of this essay, I will do two things. First, I will consider several interesting questions that apply my prior discussion to other areas of philosophy. Second, I will try to show that my conclusions are at the intuitive heart of those other questions. In this section, I will consider what my theses entail about the nature of our moral obligations generally. In §6, I will consider what my theses entail about the nature of conscious experience. In general, we can take any implication of my views to be disjunctive: either some concept picks out (directly or indirectly) a fundamental feature of the world, or else it's vague.

Consider the concept of **moral obligations**. There are some actions we're obligated to perform, some we're obligated to avoid, and some for which it's ok for us to go either way. When the reasons that ground these obligations are moral reasons, I'll call them moral obligations. Whenever an action is such that we are morally obligated to perform it – or to avoid it – I will say that our concept of **moral obligation** applies to that action. For example, the

concept of **moral obligation** clearly applies to preventing terrible suffering at no expense to me. If I can prevent a terrible murder just by saying that I would not like it, then I am obligated to do so. The concept of **moral obligation** clearly applies to that action, since I am clearly obligated to perform that action. The concept of **moral obligation** also clearly applies to causing terrible suffering for no good reason. For example, I am obligated to *avoid* committing a terrible murder just because I am bored. In this case, the concept of **moral obligation** clearly applies to committing a terrible murder out of boredom, since I am clearly obligated *not* to perform that action. On the other hand, the concept of **moral obligation** clearly does not apply to some actions. For example, I am clearly not obligated to paint my room any particular color, since everyone would be equally happy with all of my choices. No good or bad would come from my choosing blue rather than red, or orange rather than green, and so on. The concept of **moral obligation** clearly does not apply to painting my room any color.

On this description of the concept of **moral obligation**, is the concept vague? Must it always be clear whether our concept of **moral obligation** applies to an action, or could there be cases where it is not clear? In other words, could there be an action such that you are not clearly obligated to perform it, and not clearly obligated to avoid it, and not clearly allowed to go either way? One reason to think the concept of **moral obligation** must be clear is that morality is action-guiding. To see this, think of morality as a complete list of possible actions. For every action on the list, the list gives us a verdict about whether we should perform that action, avoid that action, or decide for ourselves. If our concept of **moral obligation** is vague, then this list would have some members that no one could clearly read. There would be some action on the list with no clear verdict written to one side. At best there would be hazy print,

[48]

which could provide one with only a vague sense of an answer. Does this seem right? Might

there be actions that morality delivers no clear verdict about in either direction?

I will not attempt to answer this question here. Instead, I will explore what my main

thesis and reverse thesis say about either possible answer. Suppose that we say yes; there

might be items on the list of possible actions that morality does not deliver any clear verdict

about. In other words, there might be cases in which it is vague whether the concept of **moral

obligation** applies. If the concept of **moral obligation** is vague, then my conclusions entail

several things about it. First, we can describe non-trivial application conditions for the concept

of **moral obligation** that lend themselves to a sorites series. Second, the concept must not pick

out a fundamental feature of reality. Finally, somewhere among the concept's constituents are

other vague concepts, or else it is constructed by applying vague logical operators to clear

concepts.

Alternatively, suppose that we say no; there must be a clear verdict about the moral

status of every action on the list. In other words, it is always clear whether the concept of

moral obligation applies to an action. This entails (i) that we cannot possibly give non-trivial

descriptions of the application conditions for the concept of **moral obligation**, (ii) that the

concept picks out a fundamental feature of reality, and (iii) that it is either an ultimate concept

or that it is ultimately constructed by the exclusive application of clear logical operators to

ultimate concepts.

Thus, simply by consulting our intuitions about the extent to which morality can give us

guidance, we are permitted to derive serious metaphysical conclusions. These are substantive

[49]

enough that they may sway our reactions to the other question: is it clear for every action whether we are obligated to perform it? I will now spend some time exploring these conclusions.

To begin with, it's not very plausible that our concept of **moral obligation** is itself an ultimate concept. Ultimate concepts are unanalyzable, but philosophers have been attempting to give analyses of our concept of **moral obligation** for millennia. Many of them are at least initially plausible, and they are often motivated by the reasons we offer that some action is obligatory. For example, I'm obligated to avoid senseless murder because it produces tremendous, unnecessary suffering. The plausibility of such reasons suggests that there are non-trivial application conditions for the concept of **moral obligation**. If an action would cause unnecessary suffering, then we are obligated to avoid it. Alternatively, I'm obligated to avoid senseless murder because it violates important rights. If an action would violate important rights, then we are obligated to avoid it. It is not obvious that such application conditions are trivial or metaphysically circular. If they are, it is a veiled kind of circularity – not the kind apparent in conditions like "something exists only if it is real."

If our concept of **moral obligation** is not ultimate, then it's either vague or ultimately constructed by applying clear operators to ultimate concepts. Now consider my first prediction: if the concept of **moral obligation** is not vague, then no good analysis of it will give rise to a sorites series. It is hard to say whether the concept will resist a sorites series – in principle, no matter what the correct analysis turns out to be – without simply asserting that the concept of **moral obligation** is clear. It is much easier to say of any particular analysis

[50]

whether it gives rise to a sorites series. For my purposes, I'll just try to give an example of one that does and one that doesn't. For each example, do not consider whether the proposed analysis of the concept of **moral obligation** is plausible. Instead, ask whether the possibility of a sorites series using the analysis counts for or against the proposed analysis.

First, consider an analysis of our concept of **moral obligation** that does not give rise to a sorites series. Suppose that we are morally obligated to perform an action if and only if it would produce better consequences than *any* other alternative. This description of the concept's application conditions is non-trivial. We do not need to appeal to our moral obligations in order to decide whether an action produces better consequences than any alternative. It could be that every concept used in these application conditions is itself clear and ultimate. For example, the concept of **being better than** is plausibly unanalyzable. The concepts of **consequences**, of **being an alternative**, of **being an action**, and of **being produced by** (where production is synonymous with causation) might all be unanalyzable. It would certainly be very hard to provide non-trivial application conditions for any of them. The concept of causation, for example, has proven especially resistant to substantive analysis.

If these constituents are not themselves ultimate, then it is plausible that they are constructed out of ultimate concepts and clear operators. For example, perhaps the concept of **being an alternative** is analyzable in terms of nearby possible worlds. But the concept of **being a possible world** is a very plausible candidate for an ultimate concept. Thus, on this analysis of **moral obligation**, every constituent is either ultimate itself, or constructed out of ultimate concepts and clear logical operators. Vagueness has not been introduced at any point in its

construction. It must always be clear whether the concept applies. And indeed, when we try to construct a sorites series of cases using these criteria, it is immediately obvious how it fails.

In case a_1, there are two possible actions – A and B. Action A produces consequences that are better than those produced by action B. Thus, we are clearly obligated to perform action A; the concept of **moral obligation** clearly applies to action A. In case a_2, A is still better than B, but by a slightly smaller degree. In a_3, A is still better, but by a still smaller degree. At a certain point – let's say a_{5000} – A is exactly as good as B. It is obvious that at a_{5000}, and no sooner, the concept of **moral obligation** *clearly* does not apply to A, nor does it apply to B. a_{4999} is the last case in which we have a clear moral obligation to perform A. There is no borderline region in which it is not clear. The application conditions of our concept of **moral obligation** are such that it is clear whether they are satisfied, regardless of the degree to which they are satisfied. There will always be a clear answer about what we are obligated to do.

In contrast, consider an analysis of our concept of **moral obligation** that *does* give rise to a sorites series. Suppose that we are obligated to perform some action if and only if the action does not harm any bald person. This description of the application conditions for the concept of **moral obligation** is also non-trivial. We do not appeal to the concept of **moral obligation** in order to decide whether or not an action harms any bald person. However, it is not the case that all of its constituents are clear. This analysis of the concept of **moral obligation** has the concept of **being bald** as one of its constituents. The concept of **being bald** is a paradigm vague concept. There is no doubt that this description of the concept of **moral obligation** will yield a sorites series.

In case b_1, action A would clearly harm a man with no hairs on his head. Thus, we are clearly obligated to avoid action A; the concept of **moral obligation** clearly applies to action A. In case b_2, A would clearly harm a man with one hair on his head. In b_3, A clearly harms a man with two hairs. There are a range of cases in the series for which A clearly harms a man, but it is not clear whether the man is bald. In b_{5000}, the harmed man has 4,999 hairs, such that 25% of his scalp shows through. In b_{5000} (and a neighboring range around it) although it is clear that *some* man is harmed, it is not clear whether the concept of **moral obligation** applies to A.

Suppose that you found yourself in a situation such as the one described in b_{5000}. Action A will provide a tremendous benefit to you, let's say that you will make $1,000,000 for doing A. If you are not morally obligated to avoid A, then you will definitely perform A to make the $1,000,000. However, it is vague whether you are obligated to avoid A in this case, since it is vague whether the man you will harm is bald. If someone says that you are clearly permitted to perform A, they are mistaken, since that entails that the man is clearly *not* bald. If they tell you that you are clearly obligated to avoid A, they are mistaken, since that entails that the man clearly *is* bald. Thus, there is no clear answer about what you should do; there is no clear answer about whether our concept of **moral obligation** applies to A. If baldness turns out to be a feature which determines our moral obligations, then there will be some cases for which a clear answer about our moral obligations is *impossible*. You will not be clearly guilty of doing wrong, and you will not be clearly off the hook.[20]

_____

[20] The same case serves to demonstrate that analogous normative concepts in epistemology are vague. For example, is it rational for me to harm the man for the money? The concept of

Is this plausible? Is that possibility enough for us to dismiss this proposed analysis of our

concept of **moral obligation**, independent of any other considerations? I will not say. The

purpose of this discussion is not to argue whether morality is essentially action guiding. My

purpose is to demonstrate how my criteria of vagueness can be put to work in answering these

sorts of questions. According to my conclusions, one way of deciding whether the concept of

moral obligation is vague is first to decide whether the concept permits a non-trivial analysis,

and then to decide whether that analysis gives rise to a sorites series.

Consider my second prediction: that if the concept of **moral obligation** is not vague,

then it must pick out a fundamental feature. Alternatively, if our moral obligations are

fundamental features of reality, then the concept of **moral obligation** is not vague. Is there any

good reason to think one way or the other? I have already suggested that our concept of **moral**

obligation is not likely to be an ultimate concept. Since there is plausibly *some* analysis of the

concept, it probably doesn't *directly* pick out a fundamental feature. Is there some reason to

think that it doesn't even do so *indirectly*?

If some subjective analysis of our concept of **moral obligation** turns out to be correct,

it's not likely that the concept picks out fundamental features – not even indirectly. Subjective

accounts of moral obligation appeal to the attitudes of individuals or groups. For example, you

are morally obligated to perform some action if and only if you approve of it, or if you believe it

would be good, or if it pleases you, or some such thing. Or, you are morally obligated to avoid

some action if and only if the society that you are a part of disapproves of it, or believes it

---

rational obligation will admit the same borderline cases as the concept of **moral obligation** in
the series above.

[54]

would be bad, or is offended by it.[21] Thus, on any relativist analysis of our concept of **moral**

obligation, whether the concept applies *depends* on what attitudes some set of people have

towards the action.

But it's not plausible that these are fundamental features of reality. Human beings,

their societies, and their relative beliefs and preferences don't seem fundamental at all. All of

them depend for their existence on other things, some of them even on other members of the

list. For example, human beliefs, preferences, and societies all depend for their existence on

the existence of individual humans. Individual humans depend for their existence on the

existence and behavior of micro-organisms, which depend for *their* existence on the existence

of certain molecular structures, and so on.

Maybe the fundamentality is just one more level removed. For example, suppose that

the concept of **moral obligation** can be analyzed in terms of what societies prefer. The concept

of **being preferred by society** does not itself pick out a fundamental feature, but perhaps *it* can

be analyzed in terms of more basic concepts that do. Since the relevant concepts here – such

as the concept of **being approved of by society** – are very complex, I could write at length just

analyzing *their* constituents, and their constituent's constituents, and so on.

Rather than chase the rabbit, I think it will suffice for my purposes that some of the

relevant concepts are obviously vague. For example, the concept of **being approved of** is

vague. Let's say that I approve of action *A* only if when I consider doing *A* I feel no serious

---

[21] Obviously these are toy accounts of the relativity of moral obligation. It's not relevant for my
project that the view be especially developed, but serious accounts have been developed
elsewhere. For example, see Harman and Thomson (1996), or Harman (1975).

discomfort. Although this analysis is incomplete, as much as it says seems correct. For ease of writing, I'll quantify discomfort numerically – $100d$ being very uncomfortable, $0d$ being no discomfort at all. Consider the following series. In case c_1, considering A causes me $0d$. In c_1, I also do not approve of A. In case c_2, considering A causes me $1d$. In case c_3, considering A causes me $2d$, and so on. By c_{100}, I clearly do not approve of A. What about case c_{50}? Does $50d$ clearly qualify as a serious discomfort? In case c_{50} it is vague whether I approve of A. Having no serious discomfort about X is a necessary condition for approving of X, but in case c_{50} it's vague whether I experience serious discomfort.

As a result, any analysis of the concept of **moral obligation** that appeals to the concept of **being approved of** must be vague. Any concept with a vague constituent is vague to at least that degree. But if the concept is vague, then it cannot pick out a fundamental feature – directly or indirectly. Thus, subjective analyses of the concept of **moral obligation** entail that our obligations are not a fundamental feature of reality. To the extent that one endorses a subjective analysis of **moral obligation**, my argument entails that they have reason to think that our moral obligations are not part of the fundamental structure of reality.

The point of this discussion is to show that there are extra considerations at play for philosophers who expect morality to give clear answers about what we are obligated to do. It must be that their account of **moral obligation** picks out a fundamental feature of reality, at least indirectly, which involves more careful conceptual analysis. The imagined relativist about our moral obligations cannot have morality be uniquely action-guiding. The relative concept of **moral obligation** is not ultimately built out of clear concepts and operators. If they wanted to

argue that their picture of morality really *is* uniquely action-guiding, they will have to show that each constituent of **moral obligation** is clear, and that there are no vague operators in its construction. Since I have argued that at least one of its constituents – the concept of **being approved of** – is vague, their response must show that each of *that* concepts constituents is actually clear, and so on.

I could repeat this demonstration for other normative views: for example, the view that we are obligated to perform an action if and only if God commands us to, the view that we are obligated to perform an action only if it does not disrespect other persons, and so on. I'd guess that the relevant vague constituents in those cases would be the concept of **being a command**, and the concept of **being disrespectful to persons**. If people with these views think that morality is uniquely action guiding, then they must be prepared to show that those concepts are perfectly clear.

The general point is that we can employ these strategies in a broad range of areas to determine the vagueness and fundamentality of different concepts and features. We can use our intuitions about whether morality is vague to derive conclusions about its fundamentality, and we can use our intuitions about whether morality is fundamental to derive conclusions about its vagueness.

§6 – Is Consciousness Vague?

In this section, I will consider the implications of my argument for the concept of **being conscious**. It's hard to analyze the concept of **being conscious**, but it's hard in a different way than it's hard to analyze the concept of **moral obligation**. While it's often difficult to tell whether one account of our moral obligations is better than another, it is sometimes difficult to tell whether two theories of consciousness are even about the same thing. As always, the general implication of my arguments is a disjunction: either our concept of **being conscious** is not vague, or any theory that makes consciousness a fundamental feature of reality is mistaken. If our concept of **being conscious** is clear, then any theory that does *not* make consciousness a fundamental feature of reality is mistaken.

To begin with, assume that our concept of **being conscious** is clear. I think there is good independent reason to think it is. Mostly, the natural ways of describing a case in which it is vague whether someone is conscious are not really borderline cases of **being conscious**. When I ask whether the concept of **being conscious** is vague, I mean to ask whether it can be vague whether something is having a conscious experience *at all*. Obviously it can be vague *what* one is experiencing. For example, when I have my vision tested by an optometrist, I have a hazy, fuzzy experience of the letters on her poster. This isn't really a borderline case of something's being conscious though. In that case, I am clearly conscious; it's just vague what my conscious experience is like, that is, which (non-fundamental) experiential concepts apply to it. It might be vague whether I can see the letter "E," but it is not vague whether I am having *some* conscious experience or other.

Alternatively, I may have an experience for which I have a very limited idea what its content is. For example, when I see a zebra at the zoo, I have no idea exactly how many stripes I am seeing. Unlike the optometrist case, the content of my experience is clear. There might clearly be 200 stripes on the zebra, all of which are clearly represented in my experience. Nonetheless I'm totally uncomfortable asserting that I see any particular number of stripes, because I'm just unaware of the content of my experience. Again, this is not really a borderline case of my being conscious. It is clear that I am having *some* experience, even if we are tempted to say it's vague *what* my experience is like.[22]

The case in which I am most tempted to say it's vague whether I am conscious is the case in which I am having my first or last experience. For example, sometimes I very gradually wake up from a deep sleep. There is a point at which I am clearly not having any conscious experience – while I am in deep sleep. There is a point at which I am clearly having a conscious experience – while I am fully awake. But it's not initially obvious that there's no borderline range in which it's vague whether I am having a conscious experience.[23] I do not know how to characterize such alleged borderline events, but I suspect they will turn out to be instances of one of the former "borderline" cases of **being conscious**.

Since the intuition that our concept of **being conscious** is clear is in the literature, and since the standard cases in which it is vague can be explained away, I will assume for the moment that **being conscious** is clear. I will now consider several theories of consciousness,

---

[22] The preceding discussion echoes one found in Antony (2008).
[23] If the reader is uncomfortable with questions about whether we have experiences while we sleep, we can instead think about our slow death.

broadly construed as families. On my working assumption that the concept of **being conscious** is clear, any theory that entails consciousness is not a fundamental part of reality is false. Thus, any good reason to accept a theory on which consciousness is not fundamental is a good reason to reject our intuition that the concept of **being conscious** is clear.

First, I'll consider physicalist theories of consciousness. In particular, I'll consider functionalism, type-identity theory, and eliminativism. These views all seem to provide a non-trivial analysis of the concept of **being conscious**. They at least appear to describe a dependence between consciousness and other features of the physical world. Thus, if we take this dependence seriously, these views do not count consciousness as a fundamental feature. Unless these views construct the concept of **being conscious** using only clear constituents and logical operators, this entails that the concept of **being conscious** is vague.

According to type-identity theories, the concept of **being conscious** applies to an object only if the object is in a token of a particular type of physical state. Unlike functionalist theories, type-identity theories do not define the particular types of physical state according to the role the state plays in a larger system. Rather, they define *specific* types of physical state, usually specific neural states in brains. For example, an object has narrative self-awareness only if a certain pattern of left hemisphere interpretative neural processes are occurring.

On type-identity theories, conscious states are not fundamental. There could not be conscious states without corresponding brain states. Since on this view, the concept of **being conscious** is not an ultimate concept, it's either vague, or it's constructed out of only clear concepts and logical operators. As before, I will not chase the chain of constituents to its end.

[60]

If I can show that any of the constituents of **being conscious** are vague on this view, then that will suffice.

Since the type-identity theorist almost certainly wants to analyze **being conscious** in terms of the activity of certain neural correlates, I will describe **being conscious** in terms of neurons firing. This is obviously a toy-version of very serious views in cognitive science, but any more complicated analysis can be substituted by the reader. For my purposes, this shorthand is sufficient and much simpler. Thus, suppose that an object is conscious only if it has neurons firing in a certain pattern. On this analysis, the concept of **being conscious** is vague; it gives rise to a sorites series.

Let's say that in case a_1, there is no activity in the object's relevant neurons. It is clearly not conscious. In case a_2, a positive charge has begun in the nucleus of one of the object's relevant neurons. It is still clearly not conscious. In case a_3, the charge has traveled .001 microns toward the neuron's axon.[24] The object is still clearly not conscious. In case a_{5000}, all of the object's relevant neurons are firing. The object is clearly conscious. Somewhere in the middle of the series, it is vague whether the neuron firing pattern has progressed sufficiently for the object to be conscious. Furthermore, there is nothing *besides* this firing pattern to appeal to as a type-identity theorist. The progress of the neuronal firing pattern is *all* that determines whether the object is in a conscious state. Thus, type-identity analyses of **being conscious** do not make consciousness a fundamental feature and do not build the concept out

---

[24] The average neuron of the human brain is apparently 4 microns long, measured from the furthest dendrites. For an interesting video that summarizes features of neurons, see Bosgraaf (2010).

of only clear constituents and logical operators.   In some cases it is vague whether an object is conscious.

Functionalist theories generally say that whether mental properties are present – and if so, what they are like – depends on how particular physical states function. For example, an object is in a pain state only if it's in some physical state that plays a certain functional role. This is a notion we have an intuitive grasp of. An object is a heart only if it plays the role of pumping blood in a circulatory system. Presumably, different mental properties are realized according to different functions. For example, the role that a physical state must play for the system to be a pain state is different from the role that a physical state must play for the system to be a pleasure state. Functionalist theories of consciousness in particular must specify a *general* role that a physical state must play for the system to be conscious. Thus, for example, although pain states and pleasure states are realized by different functional states, both of them realize conscious states.[25]

According to functionalist theories of consciousness, the concept of **being conscious** applies to an object only if the object is in a physical state that plays the right functional role. But if that's right, then it will sometimes be vague whether an object instantiates the relevant functional state, because it will sometimes be vague whether an object instantiates the relevant physical state. In other words, the type-identity sorites series infects the functionalist sorites series.[26] For example, a serious functionalist theory will plausibly say that whatever

---

[25] For an overview of functionalist theories of mind, see Levin (2013).
[26] A similar point is made by Antony (2008, 528): "It will then be unclear whether that brain-state bears the same causal relations (actual and counterfactual) to inputs, outputs and other

physical structures are correlated with human's being conscious instantiate whatever the

relevant functional state is. But the physical structure correlated with human's being conscious

will have to be something involving neuronal firings, such as the one described above. Since it

can be vague whether an object is in that *physical* state, it can be vague whether an object is in

the resulting *functional* state. Any hope of resolving this result with the intuition that the

concept of **being conscious** is clear will face all the same objections as the type-identity theorist

from this point on.

On eliminatavist theories of consciousness, nothing is ever conscious. Any analysis of

consciousness is too spooky, or leaves too much unexplained, or simply doesn't make sense.[27]

The result is that on these theories, the concept of **being conscious** is perfectly clear. For every

object, it is clearly not conscious. Although consciousness is not a fundamental feature on

these views, it does not permit a sorites style description of the concept of **being conscious**.

However, this is only because there is no such feature as being conscious. We will see this

again when considering dualist theories, but for now I will just say that this clearness comes at a

steep intuitive price. On its face, nothing can seem more wrong than the claim that we do not

have conscious experiences. The fact that this claim *seems* like anything at all is the best

evidence that we do.

Now I'll consider dualist views. Dualist views are the alternative to what I lumped

together as physicalist views. Dualists say that the concept of **being conscious** cannot be

---

neurophysiological states that N did, so it will be unclear whether the system realizes F. We will
thus have a borderline case of F as well. We are thus committed to this: If a property P realizes
F, then a borderline case of P is a borderline case of F."
[27] For arguments to this effect, see Ramsey (2013).

analyzed in terms of any physical concepts.[28] In particular, I'll consider Cartesian style

substance dualism, idealism, and panpsychism. It is not obvious whether these views provide

non-trivial analysis of the concept of **being conscious**. As before, unless these views construct

the concept of **being conscious** using only clear constituents and logical operators, they entail

that the concept of **being conscious** is vague.

First, consider Cartesian dualism. According to Cartesian dualists, there are two broad

types of substance in the world. There are physical substances – which include all of the

fundamental particles of physics and anything composed of them – and there are mental

substances – which include everything else. Our minds and anything else with conscious

properties are mental substances. Cartesian dualism alone is agnostic about whether mental or

physical substances are more fundamental. Cartesian dualism simply entails that mental

substances do not depend on physical substances in a metaphysically robust way. There *could*

be mental substances without physical substances, and vice versa. For all this, the fundamental

structure of reality may include features from both categories.

In this sense, Cartesian dualism doesn't give us an analysis of the concept of **being**

conscious. Really, it just tells us how we *can't* analyze the concept. It tells us that we cannot

analyze the concept of **being conscious** in terms of physical substances, since the sorts of things

that are conscious are mental substances. The latter is not reducible to the former. Substance

---

[28] I am setting aside what is sometimes called "weak" substance dualism, as it does not strike
me as a genuine kind of dualism at all. I understand this is contentious, but since my points in
this section are merely illustrative, I don't think it really matters.

dualists are likely to describe the concept of **being conscious** in terms of "what it's like," or qualia.[29]

Although it is not obvious how to interpret Cartesian dualism as an analysis of the concept of **being conscious**, I want to point out an interesting puzzle that arises from this theory. According to Cartesian dualists, is the concept of **being conscious** vague? The theory does not say that consciousness is a fundamental feature, so it seems to leave open the possibility. However, the sorts of analyses I alluded to above – those involving what-it's-likeness and qualia – are not obviously non-trivial. If a substantive analysis isn't possible, then the concept can't be vague.

Whatever account the Cartesian dualist gives, it seems like they should *somehow* connect the features of mental substances to the features of physical substances. For example, it seems that we can change what mental substances are like by changing physical substances. You could cause my mind to be in pain by stabbing my body. You can stop my mind from having conscious experiences, at least for a time, by giving my body special drugs. There is surely *some* special relationship between the two substances.

I will not try to resolve that issue here, but the Cartesian dualist faces a serious dilemma. Suppose that the Cartesian dualist builds into their analysis that what mental substances are like *depends* on physical substances. Presumably the theory will require a special non-reductive dependence. I have no idea how to characterize such a relationship, but if one can change their mental substances by stabbing their bodies, it is surely because one can

---

[29] For an overview of views like this, see §4 of Robinson (2012)

change their mental substances by affecting the neuronal structure of their brain. If so, this view seems to permit sorites style descriptions of the same kind that type-identity theories permit. In case b_1 no neurons are firing, and I clearly do not feel pain. In case b_2, the charge has been sparked in the first neuron, but I still clearly do not feel pain. And so on.

However, there is an important difference between the two sorites series. In the physicalist cases, it's crazy to say that there is some case in the series such that one is clearly conscious, but in the case before it they are clearly not conscious. That would mean that there is some point in the series at which it is vague whether the neuronal firing has taken place, but that nonetheless it is clear whether one is conscious. This is absurd, since one is conscious only if the relevant neuronal firing has taken place. My neuronal states entirely determine the quality of my conscious states. There is a strong metaphysical dependence of the former feature on the latter.

However, in the Cartesian dualist series, there is no strong metaphysical dependence. There is a mere correlation. The crazy part is not that there is some point in the series at which one sharply transitions from clearly not being conscious to clearly being conscious, even though it is vague whether their brain exhibits the relevant structure. Rather, the crazy part is that mental substances track physical ones *at all*. What do neurons firing have anything to do with an independent mind? Why is there a correspondence between the mental and physical

substances of *any* kind? Obviously there's some kind of important interaction, but if the two

substances are truly distinct, equally fundamental features, why should this be?[30]

There are at least two views that explicitly say consciousness is a fundamental feature:

panpsychism and idealism. Panpsychists say that consciousness is at the very bottom level of

reality. All things are constructed out of consciousness or mind itself. There is no adequate

analysis of the mental; indeed, the only adequate analysis of the physical world is in terms of

the mental.[31] Idealists say that what appears to be a physical world is in fact just a collection of

ideas.[32] Thus, it's not just that consciousness is a fundamental feature of reality; it is the *only*

feature of reality.[33]

Since these views take consciousness to be a fundamental feature, their concept of

being conscious must be clear. In fact, on these views the concept of **being conscious** is an

ultimate concept. Thus, the concept cannot permit a sorites series and it can have no vague

constituents. They assert that consciousness is fundamental and that the concept of

consciousness is unanalyzable. My theses rightly predict that of these views the concept is

clear.

As in §5, my aim here is not to derive any strong conclusions about the nature of

consciousness. I only intend to show that if we take the intuition that the concept of **being**

conscious is not vague seriously, then we must expect consciousness to be part of the

---

[30] This is a tired objection to substance dualism, but for a recent, concise summary of the
literature surrounding it, see Robinson (2012).
[31] For a version of this view allegedly motivated by fundamental physics, see Hameroff (1998).
[32] Most famously defended by Berkeley (1878).

fundamental furniture of reality. On popular physicalist theories of mind and consciousness, consciousness is not a fundamental feature. As a result, the concept of **being conscious** turns out to be vague. If it's implausible that the concept of **being conscious** is vague, then so much the worse for these theories. If these theories are very compelling, then we must revisit our intuition that the concept of **being conscious** is clear. On non-physicalist theories, consciousness does (or at least might) turn out to be a fundamental feature. This entails that the concept of **being conscious** is clear, but at a very high intuitive cost. These dualist views leave large explanatory gaps, and generally introduce more puzzles than they solve.

§7 – Conclusion

I began this essay with this question: are there useful ways of deciding whether a concept is vague that don't appeal to a particular theory about the nature of vagueness? In other words, can we appeal to totally general features of vagueness to decide whether a particular concept is vague? I argued that there are two criteria we can use to decide whether a concept is vague. In §1, I argued that if the concept picks out a fundamental feature, then the concept is not vague. In §3, I argued that if the concept does *not* pick out a fundamental feature, then the concept is vague. These criteria follow from perfectly general claims about how we describe concepts, what vagueness is like, and the relationship between a concept's analysis and the ontological fundamentality of what it picks out.

After defending these criteria from objections, I showed how one could use them to help answer interesting questions elsewhere in philosophy. In §5, I argued that if we assume

morality is uniquely action-guiding, then our concept of **moral obligation** cannot be vague. On that assumption, any theory of morality that does not make our moral obligations part of the fundamental level of reality, or clearly ground them in the fundamental level of reality, must be mistaken. Such theories include all relativist moral theories, theories that ground moral obligation in respect to persons, and divine command theory. If we prefer such theories, we must revisit our intuition that morality is uniquely action-guiding. In §6, I assumed that our concept of **being conscious** cannot be vague. On that assumption, any theory of consciousness that does not make consciousness part of the fundamental level of reality, or clearly ground it in the fundamental level of reality, must be mistaken. Such theories include most physicalist theories. I also showed how most dualist theories entail that the concept of **consciousness** is clear, but that they do so only at the cost of other counter-intuitive implications. These are only a few examples of how my views permit us to translate between clearness and fundamentality, and I invite the reader to explore others.

There are serious questions left open that fall outside the scope of my project. In §2, I asked whether my views are compatible with the possibility that the concept of **vagueness** is vague. In §4, I asked whether one can endorse my arguments if they accept the possibility that some supervenient concepts are unanalyzable. And although I do not address it explicitly, there is the ever present question of how vagueness is introduced into our conceptual schema *at all*. But time is a wheel, and its progress tramples all pursuits. If every claim I make is false, I am convinced that they carry at least the seeds of truth. Let greater thinkers than me replant my intellectual soil – or if it already bears fruits, let better writers harvest them on my behalf.

BIBLIOGRAPHY

Adajian, Thomas. "Moral Disagreement, Moral Realism, and Vagueness." *Applied Ethics* (2011):

1.

Antony, Michael V. "Are Our Concepts conscious state and conscious creature Vague?."

Erkenntnis 68.2 (2008): 239-263.

Antony, Michael V. "Papineau on the vagueness of phenomenal concepts."*dialectica* 60.4

(2006): 475-483.

Barnett, David. "Does vagueness exclude knowledge?." *Philosophy and Phenomenological*

Research 82.1 (2011): 22-45.

Barnett, David. "Is Vagueness Sui Generis? 1." *Australasian Journal of Philosophy* 87.1 (2009): 5-

34.

Barnett, David. "Vagueness-Related Attitudes*." *Philosophical Issues* 10.1 (2000): 302-320.

Burgess, Alexis. "A Puzzle about Identity." *Thought: A Journal of Philosophy* 1.2 (2012): 90-99.

Burgess, Alexis. "Negative Existentials in Metaphysical Debate." *Metaphilosophy* 43.3 (2012):

221-234.

Dancy, Jonathan, and Christopher Hookway. "Two conceptions of moral realism." *Proceedings*

of the Aristotelian Society, Supplementary Volumes 60 (1986): 167-205. –

Deas, Robert. "Sorensen's sorites." *Analysis* 49.1 (1989): 26-31.

Eklund, Matti. "Being metaphysically unsettled: Barnes and Williams on metaphysical

indeterminacy and vagueness." *Oxford Studies in Metaphysics* 6 (2011).

Graff, Delia. "Phenomenal continua and the sorites." *Mind* 110.440 (2001): 905-936.

Hameroff, Stuart R. "Funda-Mentality': is the conscious mind subtly linked to a basic level of the

universe?." *Trends in cognitive sciences* 2.4 (1998): 119-124.

Harman, Gilbert, and Judith Jarvis Thomson. *Moral relativism and moral objectivity*. Oxford:

Blackwell, 1996.

Harman, Gilbert. "Moral relativism defended." *The Philosophical Review* 84.1 (1975): 3-22. –

Lewis, David. "General semantics." *Synthese* 22.1 (1970): 18-67.

Perkins, Ryan. "Vagueness and the Philosophy of Perception." (2012).

Reid, Thomas. 1785/2002.Essays on the Intellectual Powers of Man, ed., Derek

Brookes.University Park: Pennsylvania State University Press.

Russell, Bertrand. "Vagueness." Australasian Journal of Philosophy, vol. 1 (1923), pp. 84-92 –

Sainsbury, Mark. "Is there higher-order vagueness?." *The Philosophical Quarterly* 41.163 (1991):

167-182.

Sainsbury, Richard Mark. "Concepts without boundaries." (1996).

Schwartz, Stephen P. 1999. "Why It Is Impossible to Be Moral*", American Philosophical

Quarterly*, 36, No. 4: 351-360.

Shafer-Landau, Russ. "Ethical disagreement, ethical objectivism and moral indeterminacy."

Philosophy and Phenomenological Research 54.2 (1994): 331-344.

Shafer-Landau, Russ. "Vagueness, borderline cases and moral realism." *American Philosophical

Quarterly* 32.1 (1995): 83-96.

Sider, Theodore. "Sorensen on unknowable obligations." *Utilitas* 7.02 (1995): 273-279.

Sorensen, Roy A. "An Argument for the Vagueness of'Vague'." *Analysis* 45.3 (1985): 134-137.

Sorensen, Roy A. "Precisification by means of vague predicates." *Notre Dame journal of formal logic* 29.2 (1988): 267-275.

Sorensen, Roy A. "Vagueness implies cognitivism." *American Philosophical Quarterly* 27.1 (1990): 1-14.

Sorensen, Roy A., and Sorensen. *Blindspots*. Oxford: Clarendon Press, 1988.

Sorensen, Roy. "Unknowable obligations." *Utilitas* 7.2 (1995): 247-71.

Tye, Michael. "Is consciousness vague or arbitrary?." *Philosophy and Phenomenological Research* 56.3 (1996): 679-685.

Tye, Michael. "Vague objects." *Mind* 99.396 (1990): 535-557.

Tye, Michael. "Why the vague need not be higher-order vague." *Mind* 103.409 (1994): 43-45.

Unger, Peter. "Conscious beings in a gradual world." *Midwest Studies in Philosophy* 12.1 (1988): 287-333

Varzi, Achille C. "Higher-order vagueness and the vagueness of'vague'." *Mind* 112.446 (2003): 295-299.

Vasile, Laura-Janett. *Moral Disagreement and Moral Indeterminacy*. Diss. Central European University, Budapest, 2010.

Wasserman, Ryan. "Personal Identity, Indeterminacy, and Obligation."

Williamson, T. 1994. Vagueness. New York: Routledge.