

THE QUIETIST'S GAMBIT

RICARDO MENA

Universidad Nacional Autónoma de México
Instituto de Investigaciones Filosóficas
rmena@filosoficas.unam.mx

SUMMARY: In this paper I develop a semantic theory of vagueness that is immune to worries regarding the use of precise mathematical tools. I call this view *semantic quietism*. This view has the advantage of being clearly compatible with the phenomenon of vagueness. The cost is that it cannot capture every robust semantic fact.

KEY WORDS: vagueness, philosophy of language, philosophy of logic, quietism, semantics

RESUMEN: En este artículo desarrollo una teoría semántica de la vaguedad que es inmune a objeciones relacionadas con el uso de herramientas matemáticas precisas. Llamo a esta postura *quietismo semántico*. Esta postura tiene la ventaja de ser claramente compatible con el fenómeno de la vaguedad. El costo que esta teoría tiene que pagar es no poder capturar todo hecho semántico robusto.

PALABRAS CLAVE: vaguedad, filosofía del lenguaje, filosofía de la lógica, quietismo, semántica

1. Introduction

It has been said that Chrysippus offered the following advice: if your chariot is heading towards a precipice, stop it well before it hits the edge. The next piece of advice has been also attributed to him: if you are confronted with a forced march sorites series, stop answering questions while you still know the answers. Both pieces of advice are worth considering seriously.¹

The second piece of advice requires some explanation. Suppose you are faced with a series of people. The first member is very tall and the last one very short. The difference in height between adjacent members of the series is only 1 mm. This is a *sorites series*

¹ I should make clear right off the bat that in this paper I want to take what seems to be a suggestion by Chrysippus and develop it in a modern framework. I am not concerned with whether the resulting view is something that fits well with Chrysippus' overall philosophy. I shall try to convey Chrysippus' views in the most theory neutral way I can. My sources of information are Barnes 1982, Bobzien 2002, and Williamson 1994.

for the predicate “ x is tall”. If you say that one of them is tall, you may want to say that the next one in the series is also tall: you do not want to imply that a millimeter can make a difference between someone who is tall and someone who is not tall. Suppose someone asks you: “Is the first one tall?”. If you feel like playing the game, you should say “Yes”. Next you are asked whether the second member is tall. Since a millimeter does not make a difference, you should answer “Yes”. Then, you get asked whether the next member is tall, and whether the one after that is tall, and so on. You’d better keep answering “Yes”, following a forced march. Also, your chariot is heading towards a precipice. If you keep answering, you should answer “Yes”. Keep doing that and you will speak falsely—you will say, of people who are not tall, that they are. You do not want to do that, so Chrysippus recommends to stop answering questions altogether. But when to stop? It is hard to tell, but if you want to make sure you are not answering falsely, you had better shut up while you still know the right answer to the question before you. That is to say, you had better stop answering at some point where you still know that, say, the n_{th} member is tall.²

There is a crucial assumption in play.³ When marching down the series, there is no way to know exactly when one should offer an answer other than “Yes”—whether it be “No”, “Neither”, “No fact of the matter”, or what have you. This, of course, is due to the fact that “Tall” is *vague*. Knowing such a thing would be to know that, say, Julia is tall, whereas Hector, who is only a millimeter shorter, has some other status. It is plausible that this is something we do not get to know. Given this, Chrysippus recommends to stop playing the forced march sorites game at some point. There is nothing you can do to assure that you answer all questions correctly—perhaps it is not even possible to do it, not even by luck.⁴ As such, either you

² For a quite different view on how to approach the forced march sorites see Shapiro 2006. As we shall see, in this paper I am not really concerned with how to approach the forced march—the focus is on how to model a vague language. More recently, Horgan (2000) offered a solution to the forced march sorites paradox that is quite similar to Chrysippus’ solution.

³ I do not wish to claim that Chrysippus endorsed the following claim. However, I do suspect that he was thinking about something along these lines.

⁴ Chrysippus’ views about this depend on whether or not he thought that there is an unknowable sharp line dividing those who are tall from those who are not tall. This issue is controversial. Williamson (1994) argues that Chrysippus did think that there are sharp boundaries and that he was in fact some kind of *epistemicist*. Bobzien (2002) challenges this interpretation. I have no views about the proper way to interpret Chrysippus. In this paper I only want to develop one of his suggestions

lose, or you stop playing. The game is broken, so it is alright to stop playing.

Let us think about a different game. For this one I want to make a substantial assumption: throughout this paper I shall assume that vagueness is a semantic phenomenon, as opposed to an epistemic or metaphysical one.⁵ Here is the game: we are faced with the task of offering an adequate semantic model for a language containing vague predicates. For simplicity I will take vague predicates to be those that can be used to set up a forced march game —i.e., “tall”, “old”, “rich”, “near downtown”.⁶ To complete our task we are given standard mathematical resources —i.e., sets, fuzzy sets, functions, objects— and a classical metalanguage. It is controversial whether or to what extent this task can be completed. Put another way, it is controversial whether this game is broken, just like the forced march game seems to be. The reasons to think that it is are quite compelling, although not decisive. If one is persuaded by them, I would recommend a Chrysippian attitude towards this game—one can play it for a while, but at some point one has to fall silent. Explaining what it means to adopt a Chrysippian attitude when offering a semantic model is one of the main objectives of this paper. First let us take a look at the reasons one may have to think that this game is broken, then we can see what the Chrysippian attitude is and why it is appealing.

Whether or not one can offer an adequate semantic model for vague languages depends on what the phenomenon of vagueness is. Let us take a look at one of the central and most puzzling features of the phenomenon. We shall proceed by way of an example.

1.1.1. Good Runners

You are observing today's 5k. First come the leading runners; they are very fast and in excellent shape. The speed and athletic excellence of the runners gradually decreases as time goes by. The runners towards the middle are not quite as fast and athletic. After some time you observe the last participants. They

in an entirely different framework. Perhaps the resulting view won't make much contact with what he actually thought.

⁵ Epistemicism and ontic theories of vagueness are serious contenders in the debate. However, this paper is about issues that concern primarily semantic theories of vagueness.

⁶ Weatherson (2010) argues convincingly that not all vague predicates are like that. In this paper I will only be concerned with predicates that are sorites susceptible. For convenience I will call them “vague predicates”.

are slow and out of shape. A friend approaches you and asks: “Did you have a chance to see good runners?” To which you reply: “Yes, all and only the fast ones were good runners.”

Here we can find a *sorites series*. The first members of the series are clearly good runners, the last members are clearly not good runners, and the running ability of the members of this series gradually decreases as they go by. Crucially, the running abilities of adjacent members of the series are indistinguishable for any practical purpose.

Now, based on your assertion, we can certainly classify some members of the series as good runners. The leading runner is clearly a good runner, given that she is very fast, and others close to her count as good runners as well. It is also clear that based on your assertion you did not classify some members as good runners; the last ones were not classified in that way. Thus, you have used “good runner” —and “fast”— to classify some members of the series in a certain way and not others. This much is uncontroversial, or at least it should be.

It is also clear that based on your assertion there is no piece of information available to us that could help us *point* at the last member of the series that has been classified as a good runner. The dominant position is that this is so because vague predicates —like “ x is a good runner” and “ x is fast”— do not draw *sharp boundaries* between cases where the predicate applies and all the rest.⁷ Relative to the sorites series in the previous example, “ x is fast” is not the kind of predicate that applies to, say, Julia but not Eli, who has been a centimeter behind Julia during the whole race.⁸ This is why this phenomenon is so puzzling. If there is no sharp boundary between the positive cases of application and all the rest, how can it be that there are both positive *and* negative cases of application of the predicate? How can it be that some of the 5k runners are fast whereas others are not? This is the kind of consideration that fuels the sorites paradox.

So here we are with our standard mathematical tool-box and a classical metalanguage trying to model this phenomenon. Let me present you with an old worry directed at this kind of model-theoretic

⁷ Of course, one has to be careful about what one means by *sharp boundary*. For now, I want to be as neutral about this as possible.

⁸ This is not a completely neutral way of phrasing the phenomenon of vagueness. An epistemicist would reject it, for instance. However, notice that the epistemicist wants to claim something similar, it is just that she cashes out the notion of *sharp boundary* in epistemic terms.

approach to the phenomenon of vagueness.⁹ The objective of the presentation is not to prove that the worry is correct —although I suspect that it is. For my purposes it is enough to show that it is reasonable and that we should take it seriously. In the rest of the paper I will argue that in the event that this worry is correct, we should adopt something like a Chrysippian attitude when modeling vague languages.

Here is the worry. Among the things we want to model is the range of application of vague predicates relative to suitable soritical domains. Much of what we say about the phenomenon of vagueness depends on this.¹⁰ Now, as we have seen, what is distinctive about vague predicates is that they do not draw sharp boundaries. What is distinctive about sets —our core mathematical tools— is that they draw sharp boundaries. A set clearly divides the objects that are members of it from all the others —a fuzzy set clearly divides the objects that are members of that set to a certain degree from all the rest. That is precisely what the range of application of a vague predicate does not do. As such, it seems odd to try to model the range of application of vague predicates by using sets —it may seem that sets are simply not the right kind of tool for this task.

There are some complications —the first one has to do with the notion of borderline case, and the second one with higher-order vagueness. A serious attempt to model vagueness as a semantic phenomenon has to be more ingenious. One may think that in between the positive and negative cases of application of the predicate there are borderline cases.¹¹ If one thinks this can be modeled with three disjoint and mutually exhaustive sets —or with two disjoint but not mutually exhaustive sets— one would face a well known objection. Vague predicates do not draw sharp boundaries between positive and borderline cases just as they do not draw those

⁹ Frege (1893) had a radical version of it, although I will focus on the version that can be found in Sainsbury 1996. Another serious proponent of this worry is Tye (1994).

¹⁰ Of course, one can also say quite a bit about the phenomenon of vagueness without a theory in the business of modeling the range of application of vague predicates. Field 2003 is a good example of this.

¹¹ I do not want to commit myself to a particular way of cashing out the notion of borderline case. One may think that a borderline case of a predicate F is an object such that it is semantically indeterminate whether the predicate applies to it, or that there is no fact of the matter as to whether it is F , or that the rules of language are silent as to whether the predicate applies to it. These are just some examples. What matters is that a borderline case is an object with a status different from that of the polar opposites.

boundaries between positive and negative cases. However, sets used in this way do draw those boundaries. Thus, the reasoning goes, if we use sets in this way we can only misrepresent the semantic values of vague predicates. It is worth pointing out that most theorists of vagueness would agree with this. However, most of them will argue that if we take into account the phenomenon of higher-order vagueness, the worry can be dismissed.

It is tempting to think that the transition from the positive to the borderline cases is vague, just as the transition from the positive to the negative cases is vague. Perhaps we can make sense of this while representing the range of application of vague predicates using sets. This line of defense typically uses a determinacy operator (D) to argue that there is no sharp division between the clear cases and the borderline cases, and the negative cases and the borderline cases. The thought is that just as there are borderline cases between the positive and the negative cases, there are second-order borderline cases between the determinate cases and the borderline cases.¹² The hope is that we can iterate this idea for each order of vagueness.

A view of this kind is certainly ingenious. It is not unreasonable to hope that this treatment of higher-order vagueness allows us to model the phenomenon of vagueness using precise mathematical tool without worrying about it. However, a worry does persist. Sainsbury (1996) makes this point in a colorful way: “This hope, however, is groundless. Indeed, its very structure should be unappealing: you do not improve a bad idea by iterating it” (p. 255). The bad idea is to use sets in order to model the range of application of vague predicates in terms of positive, negative, and borderline cases. The iteration consists in using the same idea to understand each of the orders of vagueness. If this three-fold set-theoretic distinction is too precise to model vagueness, why think that iterating it will make it better? If at all, it gives us more things to worry about. There is much room for debate here. However, even if one is optimistic about this way of understanding higher-order vagueness, it is hard not to feel the worry.

There is a further source of concern. It may be that higher-order vagueness, understood in the standard way, is paradoxical. This is a technical debate with an abundant literature. The general form of the debate is this. It is argued that given a small set of very reasonable assumptions one can prove that this way of understanding

¹² The claim that there are borderline cases at the n th order is usually formalized in the following way: $\exists x(\neg D^n F(x) \wedge \neg D\neg D^{n-1} F(x))$.

higher-order vagueness leads to inconsistencies, and therefore that it is deeply flawed.¹³ One can complain that some of the assumptions in the proofs should be rejected.¹⁴ To this one could reply that there are other proofs with similar results that do not rely on those problematic assumptions, but only on other very reasonable ones.¹⁵ And probably there can be plausible replies to those replies, and so on.¹⁶ The jury is still out. As such, one should take both sides of the debate seriously. Perhaps the standard way of thinking about higher-order vagueness is incoherent —for what it is worth, I suspect that it is.¹⁷ If it is, then we cannot appeal to it in trying to advocate for the use of standard mathematical tools in our semantic models of vague languages. If there is a justification for using those tools, it has to be quite different.¹⁸

So let me suppose that these worries are well founded, and that an attempt to use set theoretic machinery to model the *complete* range of application of vague predicates is bound to misrepresent the phenomenon one attempts to model —it is important to keep in mind this dialectical point throughout the rest of the paper. What to do then? One can declare the project futile, and abandon it altogether. I think it would be a bit soon to take this stance. Perhaps one may think that we can only aspire to offer a *homophonic* semantic theory with semantic clauses of this kind: “tall” is true of an object

¹³ Wright (1976), Williamson (1994), and Fara (2003) offer arguments to this effect. The latter two only target supervaluationist theories and others similar to them.

¹⁴ Edgington (1993), Heck (1993), McGee and McLaughlin (1995), and Asher et al. (2009) are good examples of this line of defense.

¹⁵ Fine (2008) and Zardini (2013) argue in that way.

¹⁶ Some recent attempts to get around the paradoxes are Bobzien 2013 and Pagin 2015.

¹⁷ I do not wish to deny that there is a phenomenon of higher-order vagueness. I just think that the usual way of thinking about it is problematic and that it is not clear at all that it can be used to justify the use of precise mathematical tools in our vagueness models. In section 4, I present a view that accepts the existence of higher-order vagueness. On this view, however, higher-order vagueness only plays a minimal role.

¹⁸ Cook (2002) has argued that the precision that we get from using mathematical resources should be thought of merely as an artifact of our models —an idealization, if you prefer. Cook is right in thinking that our semantic models should contain artifacts of some kind. However, it is not at all clear that when the objective is to model the phenomenon of vagueness we can afford to treat features of the model that are incompatible with that phenomenon as artifacts.

if and only if that object is tall.¹⁹ Semantic clauses of this kind are hard to reject. Also, they are not very substantial—they are trivial and not very informative.²⁰ Or perhaps one can deny that vague predicates have a range of application relative to soritical domains. The idea is that only relative to those domains are vague predicates truly paradoxical. If so, we can easily use sets to model the semantic values of vague predicates relative to non-soritical domains—these are domains where, say, there is a considerable gap between those who are tall and those who are not (for instance, everyone is either taller than 1.90 m or shorter than 1.50 m).²¹ This view is attractive, however it seems wrong to deny that “good runner” has a range of application relative to our example above—notice that in that example the domain is clearly soritical. To my mind it is part of the data to be explained that we use vague predicates to classify objects relative to soritical domains. If so, then it is not very appealing to deny that these predicates have a range of application in those cases.

In section 2, I explore a different approach. I call it *semantic quietism*. According to it, vague predicates do have a range of application relative to soritical domains, and, with a touch of Chrysippian silence, we can use sets to model that phenomenon. The worries above are up to something important, but if we are ready to sacrifice a bit, we can get around them. In section 2.3, I explain the important ways in which semantic quietism is different from other similar views in the market. In section 3, I respond to some potential worries. In section 4, I develop a version of the view that preserves classical logic for a significant fragment of the language. This is not to say that semantic quietism is committed to classical logic—only some ways of implementing it are.

2. *Semantic Quietism*

A gambit, in a chess opening, is a move that sacrifices a pawn in order to gain some positional advantage. Quietism sacrifices complete specificity; it won't offer a full representation of vague classifications. What the view gains is compatibility with the phenomenon of vagueness; this move guarantees that our representation of vague

¹⁹ Proponents of this view are Sainsbury (1996) and Ludlow (2014). Notice that I'm ignoring context dependence for simplicity.

²⁰ Although we should recognize that one can do quite a bit with them.

²¹ Rayo (2008), Gómez-Torrente (2010), and Pagin (2010) are proponents of this kind of view.

classifications does justice to the phenomenon. This is the quietist's gambit.

To a first approximation, the view I wish to explore is this: if we want to do justice of the fact that vague predicates do not draw sharp boundaries, an adequate semantic model of the range of application of vague predicates cannot capture everything that is of semantic significance. The model has to be *silent* as to whether a given predicate applies to certain objects and, in some cases, even when it is known that the predicate applies to the object. That is the semantic analog of a Chrysippian silence. Consider the predicate “ x is tall” and suppose that there are enough resources in our domain to build a sorites for it. According to quietism, an adequate semantic model for “ x is tall” is such that there is at least someone who is tall, and the model does not say anything at all about whether this person is tall—it does not even say anything at all as to whether it is indeterminate that she is tall. Of course, if our semantic model is to avoid drawing a sharp boundary between those who are not tall and the rest, there has to be someone who is not tall, and the model is completely silent as to whether she is tall.²²

2.1. The Sound of Silence

So we need to model the extension of a predicate in an incomplete way—that is what it is for the model to be silent about certain cases. Consider a simple example. Let our domain be the natural numbers from 1 to 10. The predicates “is an even number” and “is not an even number” have an extension relative to this domain. To a first approximation, here is how we can represent those extensions in an incomplete way (for convenience I will call an incomplete representation of an extension, an *incomplete extension*):

- $[[\text{is even}]] = \{2, 4, , , \}$
- $[[x \text{ is not even}]] = \{1, 3, , , \}$

This is not the official notation, but it is colorful enough to get across the main ideas—see Section 4 for a rigorous alternative. The commas with nothing in between represent the incomplete aspect of the representation. They leave open the possibility that there is something else in the extension, without committing to it. Now,

²² At this level of abstraction quietism resembles other views quite a bit. In section 3, I will make clear how this view is different in crucial respects from Soames (1999), Field (2003), and Wright (2003).

here is the sense in which these are incomplete extensions. They say something about how these predicates classify, but they do not say everything there is to say. For instance, this representation takes 2 and 4 into the positive extension of the predicate, and 1 and 3 into its negative extension, but nothing has been said about whether 5, 6, 7, 8, 9, and 10 are in the extension or counter-extension of this predicate. It is in this sense that we are dealing with an *incomplete* extensions. It is a fact that 6, 8, and 10 are in the range of application of “is even”, however this incomplete extension is entirely silent about this fact. To clarify, the representation does not say that 2 and 4 are the only members of the extension, and it does not say that 5, 6, 7, 8, 9, and 10 are negative or borderline cases. This incomplete extension is *silent* about the status of these numbers with respect of the range of application of “is even”. It is in this sense that the extension (or rather, this representation of it) is incomplete.²³

Here is an analogy that provides an intuitive way of understanding the notion of the incomplete extension of a predicate. Imagine an artist made a sculpture of you. This sculpture is very accurate in all respects, except that it is missing some parts; there is nothing corresponding to your lower torso and upper legs (see Figure on page 13). This is an *incomplete sculpture*. The sculpture on the right represents correctly the upper and lower parts of your body, but it says nothing about the middle part of your body. Now, that sculpture does not represent you as not having middle body parts, nor does it represent you as having an indefinite middle body (let us stipulate this was not the intention of the artist). Rather, that sculpture is completely silent about your middle body, and it is so in virtue of being incomplete.

Incomplete extensions, like incomplete sculptures, accurately capture some aspects of what they try to model. However, both incomplete extensions and sculptures remain *silent* about some aspects of what they represent. Incomplete extensions are silent about whether the relevant predicates apply to some members of the domain, even if there is a fact of the matter about whether the predicate applies to some of them.

²³ Perhaps incomplete extensions are of no use when modeling mathematical languages. For now I just want to elucidate the notion of an incomplete extension using a simple example. In the next subsection I will show how they can be of use in a theory of vague languages.

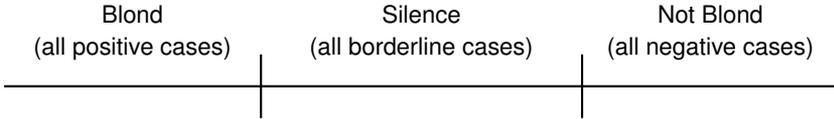
Also, notice that an incomplete extension is not some kind of metaphysically incomplete entity. Rather, it is an incomplete representation of a classical, indeterminate, or vague extension. Thanks to Andy Egan for helping me clarify this point.



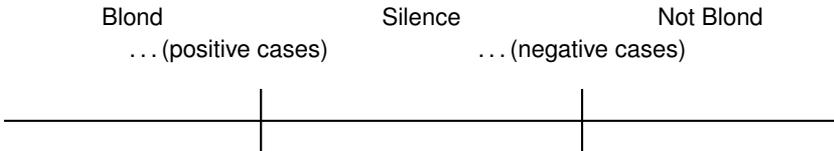
FIGURE: *Fragments*, 2016, by Bruno Catalano. Bronze Patiné, 3 pieces, 100 × 115 × 42 cm. (Photograph courtesy of De Medicis Gallery, Paris.)

2.2. Incomplete Extensions and Vagueness

Hold on. Isn't that a sharp boundary between positive cases and the silence region? Yes, but not all boundaries are harmful; it all depends on what they divide. Quietism draws harmless boundaries. A bad boundary is one that divides two distinct semantic categories, placing two objects that are indistinguishable for practical purposes on different sides of the division. This kind of boundary is bad, because what is distinctive about vague predicates is precisely that they do not draw this kind of division. A theory of vagueness that postulates that kind of cut-off fails to do justice to the phenomenon of vagueness. Here is an example of bad cut-offs (this time consider the vague predicate " x is blond"):



Now, we need to draw boundaries of some kind; that is the nature of our formal tools. We cannot draw them between all the positive cases and the first borderline case, because there is no such boundary. Where else could we draw them? According to quietism, for the boundaries to be harmless the first one should be drawn between two positive cases and the second one between two negative cases. Thus, this is the picture I recommend:



Notice that there is at least one member of the series who is blond, and our incomplete extension does not classify him as such —nor is he classified as not blond, or borderline. This little model is simply silent about the status of that member of the series with respect to the range of application of the predicate. Similarly, there is at least one member of the series that is not blond, but the model is completely silent as to whether he is blond.

Why are these boundaries not of the bad kind? The crucial thing to notice is that they are perfectly compatible with the existence of vague classifications. This incomplete way of modeling the range of application of “blond” does not misrepresent the relevant vague classification; it is just that it does not represent it completely. What makes this possible is, of course, that this incomplete extension becomes silent before it is forced to misrepresent the vague classification —this is the semantic analog of a Chrysippian attitude towards

a forced march. What a quietist boundary represents is this: it represents the point at which the semantic model becomes silent, not the transition between two distinct semantic categories.²⁴ As such, this boundary—and the one between negative cases and the silence region—is harmless.

Quietism tackles the problem of compatibility with the phenomenon of vagueness at the first-order, rather than appealing to the higher-order vagueness to get the job done. As we shall see, this puts fewer constraints on the quietist treatment of higher-order vagueness—this, in turn, opens the door to some ways of avoiding the paradoxes other theories face.²⁵

2.3. Borderline Cases

Before we move on it is important to further clarify the quietist perspective on borderline cases. As I have already pointed out, if a quietist model is silent as to whether Marco is tall, it does not thereby take Marco to be borderline tall. The silence region is not the borderline region.²⁶ As it turns out, quietism is quite flexible with respect to which conception of borderline cases one wishes to adopt. This view is perfectly compatible with a semantic conception according to which borderline cases of predicate “*F*” are those objects such that it is neither true nor false that they are *F* (Fine 1975, McGee and McLaughlin 1995, Tappenden 1993). If we go this way, we can model the range of application of a vague predicate as having positive, borderline, and negative cases, with a silence region between the positive and the borderline, and the negative and the borderline. As expected, these silence regions, if selected properly, contain some positive and negative cases or some borderline cases. In section 4, I develop a superquietist framework, showing how this may be rigorously done.

As a quietist one could also characterize borderline cases in epistemic terms (Williamson 1994). One could say, for instance, that borderline cases of predicate “*F*” are those objects such that its not knowable whether they are *F* and it is not knowable whether they are not *F*. Moreover, a quietist could think of borderline cases in terms

²⁴ Thus, the region of silence cannot be identified with the borderline region. Such an identification goes against quietism’s core views.

²⁵ As expected, the strategy will be to remain silent about certain higher-order semantic facts.

²⁶ It may be tempting to think that if a model is silent as to whether Marco is tall, it is also silent about the truth-value of any sentence universally quantifying over tall people. In section 4.5 I show that some versions of quietism do not have this consequence.

of permissibility (Shapiro 2006). For instance, she could say that an object is borderline F just in case it is permissible to assert that it is F and it is also permissible to assert that it is not F . Finally, a quietist could claim that something is borderline “ F ” just in case there is no fact of the matter about whether it is F (Field 2003). All this flexibility is possible because quietism is, first and foremost, a theory designed to ensure compatibility with the phenomenon of vagueness, not a theory designed to capture a specific conception of borderline cases.

Now, let’s contrast quietism and Crispin Wright’s agnosticism (Wright 2001 and Wright 2003). At a certain level of abstraction these two views resemble each other.²⁷ According to Wright’s view, borderline cases should be understood in terms of a particular mental state. Wright (2001) calls this mental state a *quandary*.²⁸ As such, this view recommends agnosticism regarding borderline cases and the law of excluded middle —as Wright (2003) points out, intuitionism is a natural ally. Now, and this is the point of resemblance, agnosticism and silence about p are compatible with p being true (or false). However, this similarity can only be found at the surface. Agnosticism is a propositional attitude, whereas *silence* is, primarily, a property of formal models. It is because of this that a quietist can know that Maria is tall, and, yet, let her model be silent about it.²⁹ However, one cannot know that Maria is tall and be agnostic about it —these two attitudes are flat-out inconsistent. A further difference between these two views is that agnosticism recommends that one be agnostic with respect to the law of excluded middle whereas some varieties of quietism can accept such a law (see Section 4 for more details about this). I suspect that quietism and agnosticism can play well with each other. However, it is worth keeping in mind that they exhibit sharp differences at a basic level.

It is instructive to understand the way in which quietism is different from the vagueness theory in Soames 1999. According to Soames, vague predicates are only partially defined.

²⁷ Thanks to Brian Weatherson for pointing this out.

²⁸ According to this view, a proposition p presents a quandary for subject S just in case: (a) S does not know whether or not p , (b) S does not know any way of knowing whether or not p , (c) S does not know that there is any way of knowing whether or not p , (d) S does not know that it is possible to know whether or not p , and (e) S does not know that it is impossible to know whether or not p .

²⁹ This feature of the view is what clearly distinguishes it from epistemicism (Williamson 1994, and Sorensen 2001). Of course, it is also the case that quietism is not committed to bivalence.

The first step in the construction of the model is to treat vague predicates as partially defined. As we saw [...], when F is partially defined, there will be objects to which it applies, objects to which it does not apply, and objects for which it is undefined. The latter are objects for which the linguistic rules governing the predicated and giving it meaning have nothing to say about whether or not F applies to them. (Soames 1999, pp. 206–207)

There are similarities between this view and quietism. Soames even talks about the linguistic rules governing a predicate being *silent* as to whether the predicate applies to an object. These similarities do not run deep. A quietist *may* agree with Soames in thinking that the borderline cases are those objects such that the linguistic rules are silent about them. However, the quietist insists that a semantic model compatible with the phenomenon of vagueness also has to be silent in a different way. There have to be objects such that the *model* is silent as to whether the linguistic rules governing a predicate are silent about them. That is to say, there have to be some objects that are in fact borderline cases and where the model is silent about whether or not this is the case. The quietist's silence is the one that avoids the worries stated at the beginning of this paper. In contrast, Soames' view is precisely the kind of view those worries target directly.

3. *Potential Worries*

Quietism offers an unfamiliar picture. One central aspect of this view is that, as theorists of vagueness, the best thing we can do is to build an *incomplete* model of the range of application of vague predicates. This feature of quietistic models suggests a distinctive problem. There are many different incomplete models of the very same vague language. Each model differs in how wide its silence region is — a quietist model may be silent about whether people between 1.80 m and 1.60 m are tall, whereas another may be silent with respect to people between 1.82 m and 1.63 m. The hard question is this: which quietist model *should* be selected? Or, if you prefer: which quietist model is *the best* one? These are uncomfortable questions. As soon as the quietist attempts to identify such a model, she will be stuck with horrible higher-order vagueness problems. Why is that the best model? Isn't this other very, very similar one just as good? Isn't this supposed to be vague? What is going on? Clearly, something has gone wrong.

The problem is with the question: which is *the best* quietist model? Given the puzzling features of the phenomenon of vagueness I doubt this question is coherent. It is better to think in the following terms. Some quietist models are *good enough* whereas others are not. Whether a quietistic model is good enough depends on our theoretical goals. Let's say that our objective is to model a vague language as used in a given conversation.³⁰ Relative to this goal, a quietist model is good enough if it is detailed enough to account for the communicative intentions of the participants in the conversation. Certainly, relative to this goal, a quietistic model that only classifies the tallest member of a sorites series as tall and is silent about the rest is not a particularly good model. It is not useful for our purposes. A quietist model that is complete enough to make sense of the exchange of information in the relevant conversation is good enough. Depending on the conversation, there can be many models that are good enough in this way. But, which is the best one? We should forget about that question —or at the very least keep silent when confronted with it. What matters is that we can identify and use at least one incomplete model that is good enough for our purposes.

One may keep pressing the quietist by asking her to tell us exactly which are all the quietist models that are good enough. This is another invitation to get into higher-order vagueness problems. The quietist response is the expected one. She can identify *some* of the models that are good enough, some of the models that are not good enough, and then remain silent about the rest. Her answer would be incomplete, but it will also guarantee that it is not misrepresenting the vagueness of “*x* is a good enough model for purpose *y*”.

Here is a related worry. Suppose we select a good-enough quietist model. Let's call it *M*. Right away we can identify a slightly better one. Here is why. If *M* is a good quietist model, then there is a member of the domain who is, say, tall, and that the model is silent about it —we can call such a member “Stiglitz”. If so, there has to be a good-enough quietist model *M'* such that according to it Stiglitz is tall. *M'* is better than *M* simply because it captures one more semantic fact. Now, if we select *M'* rather than *M*, we can be sure there will be an *M''* that captures one more semantic fact. So we better select *M''*. Once again our chariot is heading towards the precipice. As before, the quietist recommends to stop it well before it hits the

³⁰ If one attempts to model a language in the abstract —in isolation from any instance of use— the quietist will have trouble finding constraints for the selection of good models. I do not see how this could become a solid objection.

edge. The quietist has to be content with a good-enough model, even if she knows there is a better one available. The payoff is that she can rest assured that her model is compatible with the phenomenon of vagueness. This is another aspect of the quietist's gambit.

There is another objection the quietist has to respond to. It can be phrased as follows: quietism is not a theory of vagueness; it is, in fact, a non-theory of vagueness. In a sense this is correct. Theorists of vagueness —myself included— dream of a theory of vague classifications that illuminates the mysterious transition from the positive to the negative cases. Quietism is not the theory that we have been dreaming about. However, keep in mind that if the worries at the beginning of this paper are correct, it is not at all clear that we can have a better theory. As such, the quietist should not mind this particular objection —it is not the quietist's fault, it is just that the phenomenon of vagueness does not allow for more.

4. *Semantics, Logic, and Paradox*

Quietism is not committed to a particular logic, and there are many technical details regarding a semantic theory that this view is neutral about. Which semantics and logic we build on top of its framework have a major influence on how we approach central topics in the philosophy of vagueness, language, and logic. For instance, which solution to the sorites paradox we endorse, as well as our treatment of penumbral connections, standards of comparison, forced march sorites, the interaction between context shifts and vague expressions, among others, depend on this. Quietism, without further supplementation, is neutral about these issues. The only thing that this view promises is that a good implementation of it will be compatible with the phenomenon of vagueness. To the best of my knowledge this is the only theory that can promise such a thing with a straight face, while still relying on standard mathematical tools.

In the next section I develop superquietism. This view is the fusion of quietism and supervaluationism. Superquietism, unlike supervaluationism, is clearly compatible with the phenomenon of vagueness; there are no nasty sharp boundaries to worry about. Also, this view preserves classical logic for the *D*-free language fragment. The sorites paradox gets treated in the standard supervaluationist way.

4.1. Superquietism

The core superquietist idea is that it is problematic to understand (super)truth in terms of quantification over *all* admissible (classi-

cal) interpretations of the language —to do so is a good way of getting into higher-order vagueness problems. Instead, what the theory does is to understand (super)truth in terms of quantification over *some* admissible interpretations and a region of silence comprising some admissible and some inadmissible interpretations. As technical details will soon reveal, by going this way superquietism can think about vague classifications as follows: relative to a sorites series for “tall”, there is a silence region between the positive and the borderline cases, and the negative and the borderline cases. These regions of silence are just what you would expect; they contain some positive cases, some negative cases, some borderline cases, and a paradoxical region we had better be silent about.

The superquietist approach to the sorites paradox is, to some extent, supervaluationist. In both theories (super)truth is a matter of being true in a certain collection of classical interpretations —the difference lies in what they take this collection to be. Given that in both cases the interpretations are classical, the inductive premise of the sorites paradox is always false, and its negation is always true. As is explained in the next section, superquietism also validates all penumbral connections formulated in the object language.

4.2. Admissible Interpretations

Our starting point is the notion of an *admissible interpretation*. In a standard supervaluationist framework (Fine 1975, McGee and McLaughlin 1995, Keefe 2000) the range of admissible interpretations is what determines, in good part, the division between positive, borderline, and negative cases. Thus, if we want to include a region of silence between the positive and the borderline, and between the negative and the borderline, we have to play around with the range of admissible interpretations. Given that “ x is an admissible interpretation” is vague by any reasonable standards, the quietist should use her tools to interpret this predicate. Here is one way in which this can be done. We can represent the positive extension of “ x is an admissible interpretation” as a pair $\langle P_A, S_A \rangle$ and the negative extension as $\langle N_A, S_A \rangle$, where P_A contains *some* admissible interpretations, N_A contains *some* inadmissible interpretations, and S_A comprises the silence range —where the subscript “ A ” stands for “admissible”.³¹ Needless to say, S_A should contain a few admissible and a few in-

³¹ It is worth mentioning that, strictly speaking, S_A isn’t needed in the formalism. After all, this set contains all the interpretations that are not in P_A and N_A , so everything we can express with S_A we can also express without it. The only reason

admissible interpretations. This, again, is how quietism gets around higher-order vagueness worries.

How should we decide which interpretations go into P_A and S_A ? The answer to this question heavily relies on our particular theoretical goals. Normally we would be attending to a particular conversation and we would have a set of goals we want our model to meet. For the sake of simplicity, let's abstract away from these complications and focus on a language whose only predicate is "tall" and whose domain of discourse only contains a sorites series for this predicate. Let us call this series "T" and let's assume that it has 100 members, where 1 is the shortest and 100 is the tallest. (Clearly, a sorites for this predicate would have more members. I am just simplifying the example in a harmless way). An admissible interpretation of this language—classical as it is—draws a sharp boundary somewhere in T. If an interpretation draws a boundary between, say, 41 and 42, then, on this interpretation all the members of the series between 1 and 41 (inclusive) are not tall, and the rest are tall. If two interpretations are different, then they classify the members of T differently—they draw the cut-off at different locations in T.³² A good way of selecting the members of P_A is by taking a cluster of interpretations that draw cut-offs in a sub-region of the borderline region—these are uncontroversial admissible interpretations. Let's assume, for simplicity's sake, that this sub-region contains the members 40–60. Notice that this is only a proper sub-region of the borderline region, meaning that, at the very least, 39 and 61 are also borderline cases. Thus, the interpretations in P_A are the ones that draw a cut-off at any place between 40 and 60.

Now, there are interpretations that draw the cut-off to the left (towards 1) and others to the right (towards 100) of this borderline sub-region. Some of them go into S_A . How many? For sure the one that draws the cut-off between 39 and 40, and the one that does it between 60 and 61; these are some of the admissible interpretations that the model is *silent* about. Which others? Well, we have to go far enough as to include some inadmissible interpretations that draw the cut-off among clearly tall/not-tall members of T. Recall that a good region of silence must include some of the clear cases (clearly admissible interpretations in this case) and some clear negative cases (clearly inadmissible interpretations). Let's say that 20 and 19 are

I include S_A is because I find it easier to explain the semantics this way. Thanks to Agustín Rayo for helpful comments here.

³² I'm assuming that these interpretations satisfy all penumbral connections.

clearly not tall. Let's also say that 79 and 80 are clearly tall. An interpretation that draws a cut-off between any of these pairs is inadmissible. We can say, then, that the members of S_A are the interpretations that draw a cut-off somewhere between 40 and 20, or somewhere between 60 and 80.³³ All the interpretations that are not in P_A or S_A , are in N_A .

It is important to notice that our selection of P_A , S_A , and N_A was *semi-arbitrary*. We could have chosen a slightly bigger or smaller S_A , and, as a consequence, a bigger or smaller P_A and N_A . Whether our choice was a good one is determined by how well we can capture the flow of the particular conversation we are concerned with. Our choice, however, was not completely arbitrary; no inadmissible interpretation could have been admitted in P_A and no admissible one could have been admitted in N_A .

4.3. Supertruth

We have taken our first step—we can understand the range of admissible interpretations in a way acceptable by quietist standards. The second step is to find a good way of defining *supertruth* and the *determinacy* operator. Given that this operator will be defined in terms of supertruth, we shall focus on this notion first.

Textbook supervaluationism defines supertruth as truth in all admissible interpretations. It may seem, then, that superquietism should define supertruth as true in all the interpretations in P_A . Doing so, however, leads to disaster. Consider 60 in our previous example. She is borderline tall, but the model is silent about it. However, “60 is tall” is true in all the interpretations in P_A . This is so because all the interpretations in this set draw the cut-off somewhere between 40 and 60, so there is no interpretation in P_A that classifies 60 as not tall. Therefore, if supertruth is true in all the interpretations in P_A , then “60 is tall” is supertrue. But this is wrong, since 60 is borderline tall. A moment of reflection shows that a version of this result obtains in every proper superquietist model.

The way around this problem is to let supertruth quantify over the union of P_A and S_A . This is perfectly in line with quietism's ideology; given that “admissible interpretation” is vague, it is not clear that we

³³ Thus, for example, the interpretation that classifies all the members between 1 and 25 as not tall, and the rest as tall, is in S_A . Also, the interpretation that classifies the members between 1 and 67 as not tall and the rest as tall is in S_A . However, the interpretation that classifies the members between 1 and 5 as not tall and the rest as tall is *not* in S_A (it is in N_A).

can quantify over all and only the admissible interpretations without drawing nasty boundaries. Perhaps the best thing we can do is to quantify over *some* admissible interpretations and a region of silence. This may sound risky, because in S_A there are some interpretations that are not admissible. However, if we are careful enough in how we assign supertruth-conditions, no unpleasant result will follow. In the next section I will show how to do that.

4.4. Semantic Clauses and the Silence Region

Before we cash out these ideas with more precision, let's settle on some terminology. By an admissible interpretation of the language we shall understand —following McGee and McLaughlin 1995— a classical model that satisfies both *penumbral connections* and *classificatory constraints*. Each of these models is a complete *sharpening* of the *base model*.³⁴ The base model is a non-classical model of vague language with a partial interpretation function —it is easy to interpret a partial function in a quietist way. Following quietism, we represent the positive extension of “ x is an admissible interpretation of L ” as the pair $\langle P_A, S_A \rangle$ —recall that if we choose this pair wisely, there are some admissible and some inadmissible interpretations in S_A .

Given this, we partially define *supertruth* in a model and the determinacy operator (D) as follows (where Φ is a formula, M is a model, and σ is a variable assignment):³⁵

1. Φ is *supertrue* in M, σ , if for all $M' \in P_A \cup S_A$, Φ is true in M', σ .
2. Φ is *not supertrue* in M, σ , if there is a $M' \in P_A$, such that Φ is false in M', σ .
3. Φ is *superfalse* in M, σ , if for all $M' \in P_A \cup S_A$, Φ is false in M', σ .
4. Φ is *not superfalse* in M, σ , if there is a $M' \in P_A$, such that Φ is true in M', σ .
5. $D\Phi$ is true in M, σ , iff Φ is supertrue in M, σ .

³⁴ Shapiro 2006 offers good reasons to consider partial sharpenings as well. However, given that we are abstracting away contextual factors we can safely ignore his arguments.

³⁵ These definitions are partial because we don't offer conditions that are both necessary and sufficient, except in clause 5.

There are two very important things to notice right away. The first one is that in clauses 1–4 we don't offer conditions that are *both* necessary and sufficient —necessary and sufficient conditions typically repel vagueness. The second one is that, in 1 and 3, both P_A and S_A are built into the sufficient condition, but in clauses 2 and 4 only P_A appears in the sufficient condition. Notice, however, that these sufficient conditions are quite uncontroversial and, as it will become clear soon, this asymmetry in the clauses is exactly what we need.

These clauses give us precisely what we wanted —a region of silence between the positive, borderline, and negative cases. This is, of course, only provided that we choose P_A , S_A , and N_A wisely. For an object to be borderline tall, it has to satisfy $\neg DTall(x) \& \neg D\neg Tall(x)$. For this to be the case it is enough if a person is tall relative to a model in P_A and not tall relative to a different model in P_A . For a superquietist model to place someone in the silence region between the positive and borderline cases, that person has to be tall in all the models in P_A and not tall in some models in S_A . A model like this is silent as to whether it is supertrue that that person is tall, since the sufficient conditions in clauses (1) and (2) are not satisfied. Therefore, the model is also silent as to whether it is determinate that she is tall. Hence, it is also silent as to whether she is borderline tall. This can be so, of course, even if it is in fact determinate that that person is tall —for this to be the case, she only has to be not tall in those models in S_A that are inadmissible. A good superquietist model has cases like that. It is trivial to specify what it would take for a superquietist model to place someone in the silence region between the borderline and negative cases.

4.5. Superquietist Logic

It is time to say something about the notion of *validity*. Traditionally there are at least two options open to Supervaluationism; global and local validity.³⁶

- Global Validity: $\Gamma \models_g \Phi$ iff every $\Psi \in \Gamma$ is true in all admissible interpretations, then Φ is true in all admissible interpretations.
- Local Validity: $\Gamma \models_l \Phi$ iff for every admissible interpretation, if every $\Psi \in \Gamma$ is true in it, then Φ is true in it.

Thus, a globally valid argument is one that preserves supertruth, and a locally valid argument is one that preserves truth in all admissible interpretations.

³⁶ For more alternatives, see Varzi 2007.

Whether supervaluationism should accept global or local validity has been the subject of debate. Williamson (1994) and Fara (2010) argue that supervaluationism must hold to global validity and that doing so leads to horrible problems. Keefe (2000) argues that even though supervaluationism should accept global validity, there are no problems that emerge from doing so. McGee and McLaughlin (1995) and Varzi (2007) argue that the right notion of validity is the local one.

In this paper I shall only be concerned with what superquietism looks like if we accept local validity. Now, it is clear that we need to modify the definition of local validity if it is to be acceptable by quietist standards. In particular, in our definition we have to allow for quantification over all the interpretations in $P_A \cup S_A$. The resulting clause is this:

- $\Gamma \models_l \Phi$ if for every $M \in P_A \cup S_A$, if every $\Psi \in \Gamma$ is true in M , then Φ is also true in M .

The superquietist version of local validity delivers some nice results. Given that all the interpretations in $P_A \cup S_A$ —for whatever choice of P_A and S_A —are classical, all classical valid inferences hold in them. Therefore, all classically valid inferences hold according to this version of superquietism. Also notice that all instances of $\Phi \vee \neg\Phi$ and any other logically valid sentence are valid (and supertrue) given that they are true in all the (classical) interpretations in $P_A \cup S_A$.³⁷ Please keep in mind that this is *not* to say that quietism assumes classical logic—it is only this particular way of developing the view that has this assumption. There can be non-classical ways of being a quietist.

Another nice feature is that all D -free penumbral connections are valid. The reason for this is that even though not all interpretations in $P_A \cup S_A$ are admissible—because not all interpretations in S_A are—all of them satisfy penumbral connections. Some interpretations in S_A are not admissible because they violate some classificatory constraints—they classify clearly tall/not tall people as not tall/tall—however, that is not to say that they violate penumbral connections. Given that all penumbral connections are true in all interpretations in $P_A \cup S_A$, all penumbral connections are supertrue.

³⁷ We can only say this for the fragment of the language that is D free. Once we introduce this operator, we have to remain silent about whether certain sentences are valid. More about this below.

Let's consider a specific example. Julia is slightly taller than Carla. Thus, the following sentence is an instance of a penumbral connection:

(a) If Carla is tall, then Julia is tall.

As such, (a) is supertrue according to superquietism. This can be shown simply by pointing out that all interpretations in $P_A \cup S_A$ respect all penumbral connections. Hence, (a) is true in all of them. The very same thing can be said about every other penumbral connection —so long as the sentence in question does not contain D .

As it turns out, superquietism can tackle the sorites paradox in the same way supervaluationism does.³⁸ According to superquietism, relative to a sorites series for F , (b) and (c) are both supertrue (where x' is the successor of x):

(b) $\neg\forall x(F(x) \rightarrow F(x'))$

(c) $\exists x(F(x) \wedge \neg F(x'))$

This is so because, regardless of our choice of P_A and S_A , all the interpretations in these two sets are classical and satisfy all the penumbral connections. Thus (b) and (c) are true in all the interpretations in these sets.

As expected, relative to a soritical domain for predicate F , the following sentence is supertrue:

(d) $\neg\exists xD(F(x) \wedge \neg F(x'))$

Relative to this kind of domain, this sentence is always supertrue according to superquietism. Notice that a good superquietist model won't be silent as to whether that sentence is supertrue.³⁹ Consider an arbitrary member of the domain, a . Suppose the model is silent as to whether $DF(a)$. Then, a is in the silence region between the positive and borderline cases.⁴⁰ Thus a' is in that same silence region or it is a borderline case —this, of course, is an inclusive disjunction.

³⁸ There are new developments in logic that have a more sophisticated approach to the paradox. Good examples of this are Cobreros et al. 2012, and Zardini 2014. Quietism should be friendly to this kind of approach. However, this is something that remains to be shown.

³⁹ I shall assume, as seems reasonable, that if a model is silent as to whether p , and in that model $\neg q$ is supertrue, then $\neg(p \wedge q)$ is supertrue.

⁴⁰ It is easy to see that a couldn't be in the silence region between borderline and negative cases if the model is silent as to whether $DF(a)$.

If the semantic model declares it a borderline case, then it follows right away that $\neg D\neg F(a')$. If a' is in the silence region between the positive and the borderline cases, then $F(a')$ is true in all the members of P_A and false in at least one member of S_A . But then it is not supertrue that $\neg F(a')$ since that sentence is false in at least one model in P_A (in fact it is false in all of them). Either way, it is not supertrue that $\neg F(a')$ and, therefore, it is the case that $\neg D\neg F(a')$. Now suppose that the model is silent as to whether $D\neg F(a')$. Then, a' is in the silence region between the borderline and negative cases. Given that a is the predecessor of a' , it is in that same silence region or it is a borderline case. Either way, it is not supertrue that $F(a)$. Therefore, it is the case that $\neg DF(a)$. Hence, if the model is silent about one conjunct, the other one is false. Thus, we can conclude that $\neg(DF(a) \wedge D\neg F(a'))$. Since a is arbitrary we can conclude that (d) is the case.

We must be a bit careful here. One may argue that (d) entails a first-order gap principle: $\forall x(DF(x) \rightarrow \neg D\neg F(x'))$. If higher-order gap principles are true in a superquietist model, along with suitable iterations of D in front of them, then we can run into the paradox presented in Zardini 2013. In order to avoid this problem the superquietist can remain *silent* as to whether some of those higher-order gap principles are true. In fact, in order to dodge the paradox it is enough to remain silent about whether some of those principles are true once we have iterated D in front of them a few times. This is easy to do by including a non-transitive accessibility relation into our semantics in the standard way. Given such an accessibility relation, different models may have a different $P_A \cup S_A$ assigned to them.

So let's describe a case where $D\Phi$ is true in a model M and the same model is silent about whether $DD\Phi$ is true. Given that $D\Phi$ is true in M , Φ is true in every model in $P_A \cup S_A$ relative to M . Now, let's make M be silent about $DD\Phi$ by having a M' in S_A that is silent about $D\Phi$. (If M' is silent about $D\Phi$, then we cannot say whether $D\Phi$ is supertrue relative to M and thus we have to be silent as to whether $DD\Phi$ is true relative to M .) Thus, we get that Φ is true in all models in P_A relative to M' and false in at least one model in S_A relative to M' —from this it follows that S_A is not the same relative to M and M' . So, when Φ is a gap principle we can specify a situation where $D\Phi$ is true in a model, even though that same model is silent about whether $DD\Phi$.

Superquietism is not as pretty as supervaluationism; it remains silent about sentences supervaluationism declares valid. This is the

price superquietism pays in order to guarantee compatibility with the phenomenon of vagueness. Here is a concrete example. In a super-valuationist framework the following sentence is always supertrue:

$$(i) \quad DF(a) \vee \neg DF(a)$$

However, sometimes, in a superquietist framework we have to be silent as to whether sentences of this form are true —which is *not* to say that they are false or indeterminate. Consider 6l in our example above. Our model is *silent* as to whether $DTall(6l)$ —6l is tall in all the interpretations in P_A , but not in all the interpretations in S_A . The model is also silent as to whether $\neg DTall(6l)$ —there is no interpretation in P_A where 6l is not tall. Therefore, the model is silent as to whether $DTall(6l) \vee \neg DTall(6l)$. Similarly, our model is also silent about $DTall(6l) \rightarrow DTall(6l)$ for the exact same reasons.

A superquietist model can be silent about pretty much any logical truth construed using D . Three things are worth noting in order to make this fact a bit more palatable. First, when a model is not silent about logical truth A , then A comes out as supertrue. Second, superquietism has no problem validating all the inference we would want a nice logic to validate. So, for example, even though a model can be silent about $DF(a) \vee \neg DF(a)$, $DF(a) \rightarrow F(a)$, and $DF(a) \rightarrow (F(a) \vee G(b))$, the following inferences are always valid: $DF(a) \vDash_l DF(a)$, $DF(a) \vDash_l F(a)$, and $DF(a) \vDash_l (F(a) \vee G(b))$.⁴¹ This, of course, is a consequence of the fact that all models in $P_A \cup S_A$ are classical and that we are using local validity. Third, if one thinks that D is in the metalanguage, as McGee and McLaughlin (1995) do, then it is a bit less worrisome that a model can be silent about logical truths construed using that operator —this way we can rest content that our object language is well-behaved.

It is worth noting that a superquietist model can also be silent about instances of penumbral connections involving D in the silence range. For example, given that our model is silent as to whether $DTall(60)$ and $DTall(61)$, it is also silent as to whether $DTall(60) \rightarrow DTall(61)$. To my mind this is an acceptable price to pay. Superquietism is compatible with the phenomenon of vagueness and preserves classical logic for the fragment of the language that is D -free. These are nice features that no other theory can claim in an uncontroversial way. It is true that the fragment of the language that contains D is not as well-behaved as we may wish. However, keep in mind that if D is only a metalinguistic operator, we should not worry

⁴¹ From this it follows that conditional proof doesn't hold in this logic.

that much, so long as we get the penumbral connections right in the object language. That we can do.⁴²

REFERENCES

- Asher, N., J. Dever, and C. Pappas, 2009, "Supervaluations Debugged", *Mind*, vol. 118, no. 472, pp. 901–933.
- Barnes, J., 1982, "Medicine, Experience and Logic", in J. Barnes, J. Brunschwig, M. Burnyeat, and M. Schofield (eds.), *Science and Speculation*, Cambridge University Press, Cambridge, pp. 24–68.
- Beall, J.C. (ed.), 2003, *Liars and Heaps: New Essays on Paradox*, Clarendon Press, Oxford.
- Bobzien, S., 2013, "Higher-Order Vagueness and Borderline Nestings: A Persistent Confusion", *Analytic Philosophy*, vol. 54, no. 1, pp. 1–43.
- , 2002, "Chrysippus and the Epistemic Theory of Vagueness", *Proceedings of the Aristotelian Society*, vol. 102, pp. 217–238.
- Cobrerros, P., P. Egge, D. Ripley, and R. van Rooij, 2012, "Tolerant, Classical, Strict", *Journal of Philosophical Logic*, vol. 2, no. 41, pp. 347–385.
- Cook, R.T., 2002, "Vagueness and Mathematical Precision", *Mind*, vol. 111, no. 442, pp. 225–248.
- Dietz, R., and S. Moruzzi (eds.), 2010, *Cuts and Clouds: Vagueness, its Nature and its Logic*, Oxford University Press, Oxford.
- Edgington, D., 1993, "Wright and Sainsbury on Higher-Order Vagueness", *Analysis*, vol. 53, no. 4, pp. 193–200.
- Fara, D.G., 2010, "Truth in a Region", *Vagueness and Language Use*, Palgrave Macmillan, Oxford, pp. 222–248.
- , 2003, "Gap Principles, Penumbral Consequence, and Infinitely Higher-Order Vagueness", in Beall 2003, pp. 195–221.
- Field, H., 2003, "No Fact of the Matter", *Australasian Journal of Philosophy*, vol. 81, no. 4, pp. 457–480.
- Fine, K., 2008, "The Impossibility of Vagueness", *Philosophical Perspectives*, vol. 22, no. 1, pp. 111–136.
- , 1975, "Vagueness, Truth and Logic", *Synthese*, vol. 30, no. 3, pp. 265–300.
- Frege, G., 1893, *Grundgesetze der Arithmetik: Begriffsschriftlich abgeleitet*, vol. 1. H. Pohle, Jena.
- Gómez-Torrente, M., 2010, "The Sorites, Linguistic Preconceptions, and the Dual Picture of Vagueness", in Dietz and Moruzzi 2010, pp. 228–253.
- Heck, R.G., 1993, "A Note on the Logic of (Higher-Order) Vagueness", *Analysis*, vol. 53, no. 4, pp. 201–208.

⁴² For their many helpful comments, I'm grateful to Brian Weatherson, Agustín Rayo, Brian McLaughlin, Ernie Lepore, Lisa Miracchi, Andy Egan, Axel Barceló, Will Fleisher, and an anonymous referee.

- Horgan, T., 2000, "Facing Up to the Sorites Paradox", in *The Proceedings of the Twentieth World Congress of Philosophy*, vol. 6, pp. 99–111.
- Keefe, R., 2000, *Theories of vagueness*, Cambridge University Press, Cambridge.
- Ludlow, P., 2014, *Living Words: Meaning Underdetermination and the Dynamic Lexicon*, Oxford University Press, Oxford.
- McGee, V. and B. McLaughlin, 1995, "Distinctions without a Difference", *The Southern Journal of Philosophy*, vol. 33, no. 1, pp. 203–251.
- Pagin, P., 2015, "Tolerance and Higher-Order Vagueness", *Synthese*, pp. 1–34.
- , 2010, "Vagueness and Central Gaps", in Dietz and Moruzzi 2010, pp. 254–272.
- Rayo, A., 2008, "Vague Representation", *Mind*, vol. 117, no. 466, pp. 329–373.
- Sainsbury, R.M., 1990, "Concepts without Boundaries", inaugural lecture delivered at King's College London, Department of Philosophy, November, 1990.
- Shapiro, S., 2006, *Vagueness in Context*, Oxford University Press, Oxford.
- Soames, S., 1999, *Understanding Truth*, Oxford University Press, Oxford.
- Sorensen, R., 2001, *Vagueness and Contradiction*, Oxford University Press, Oxford.
- Tappenden, J., 1993, "The Liar and Sorites Paradoxes: Toward a Unified Treatment", *The Journal of Philosophy*, vol. 90, no. 11, pp. 551–577.
- Tye, M., 1994, "Sorites Paradoxes and the Semantics of Vagueness", *Philosophical Perspectives*, vol. 8, pp. 189–206.
- Varzi, A.C., 2007, "Supervaluationism and its Logics", *Mind*, vol. 116, no. 463, pp. 633–676.
- Weatherson, B., 2010, "Vagueness as Indeterminacy", in Dietz and Moruzzi 2010, pp. 77–90.
- Williamson, T., 1994, *Vagueness*, Routledge, London.
- Wright, C., 2003, "Vagueness: A Fifth Column Approach", in Beall 2003, pp. 84–105.
- , 2001, "On Being in a Quandary: Relativism Vagueness Logical Revisionism", *Mind*, vol. 110, no. 437, pp. 45–98.
- , 1976, "Language-Mastery and the Sorites Paradox", in G. Evans and J. McDowell (eds.), *Truth and Meaning: Essays in Semantics*, Clarendon Press, Oxford, pp. 223–247.
- Zardini, E., 2013, "Higher-Order Sorites Paradox", *Journal of Philosophical Logic*, vol. 42 no. 1, pp. 25–48.
- , forthcoming, "First-Order Tolerant Logics", *Review of Symbolic Logic*.

Received: September 20, 2016; revised: May 16, 2017; accepted: August 7, 2017.