SEMANTIC QUIETISM

BY RICARDO MENA-GALLARDO

A dissertation submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Philosophy

Written under the direction of

Brian McLaughlin

and approved by

New Brunswick, New Jersey

October, 2014

ABSTRACT OF THE DISSERTATION

Semantic Quietism

by Ricardo Mena-Gallardo

Dissertation Director: Brian McLaughlin

In *Semantic Quietism*, I defend and develop a view according to which only semantic theories that are incomplete in important respects can be compatible with the phenomenon of vagueness. I defend this view by arguing that any semantic theory that attempts to capture every robust semantic fact is incompatible with this phenomenon. Then I show how a semantic theory that is incomplete in the right way can have this kind of compatibility. I call this view Semantic Quietism. Finally, I show how this view can be used to solve some problems had by theories of linguistic communication.

Acknowledgements

Many thanks to the members of my dissertation committee, Brian McLaughlin, Ernest Lepore, Andy Egan, and Agustín Rayo. I'm specially grateful to Brian Weatherson for being extremely patient and helpful when my project was in its early stages.

Thanks also to all my friends, colleges, and professors that influenced my work in one way or another: Andy Egan, Alex Anthony, Axel Barcelo, Ben Levinstein, Bob Beddor, Bruno Velazquez, Carla Merino, Carlotta Pavese, Carlos Montemayor, David Black, Dean Zimmerman, Dirk Kindermann, Eduardo García-Ramírez, Eliot Michaelson, Gabe Greenberg, Heather Demarest, Jason Stanley, Jeff King, Jeniffer Wang, Jonathan Schaffer, Kurt Sylvan, Lisa Miracchi, Lourdes Valdivia, Lucy Jordan, Nick Beckstead, Nick Tourville, Marco Dees, Maite Ezcurdia, Megan Feeney, Meghan Sullivan, Michael Hicks, Michael Johnson, Moisés Macías, Owen, Peter van Elswyk, Rodrigo Borges, Simon Goldstein, Stephanie Leary, Ted Sider, Thomas Blanchard, Tim Campbell, Tobias Wilsch, Tom Donaldson, Tyrion Lannister, Una Stojnic, Valeria Luiselli, Van McGee, Will Fleisher, Zachary Miller, and many people that I've forgotten to mention.

For their unconditional support, thanks to Sofía Gallardo, Armando Mena, Alejandra Mena, Andrea Gallardo, and Christian Staub. Thanks to Lisa Miracchi for helping me escape darkness and making me incredibly happy.

Dedication

To Sofía Gallardo; sometimes beauty does not belong to this world.

Table of Contents

Abstract ii

Acknowledgements iii

Dedication iv

1. Introduction 1

2. On Cutting-Off Vagueness 7

2.1. An Adequacy Condition 7

2.2. Exhaustive Theories of Vagueness 11

2.3. Cut-Offs, Cut-Offs Everywhere 13

2.4. Higher Problems 15

2.5. Vagueness Must Stop 19

2.6. Instruments, Artifacts, and Tools 25

3. The Quietist's Gambit 32

3.1. The Phenomenon 32

3.2. Quietism 34

3.3. Frege's Paradise 36

3.4. Fall from Paradise 38

3.5. The Sound of Silence 41

3.5.1. Incomplete Specifications and Vagueness 44

3.5.2. Borderline Cases 47

3.6. Potential Worries 49

3.7. Semantics, Logic, and Paradox 52

3.8. Superquietism 54

4. Silence in Conversation . . . 65

4.1. Introduction . . . 65

4.2. The Stalnakerian Framework . . . 66

4.3. Pragmatic Evasion . . . 68

4.3.1. The Accommodation Strategy . . . 72

4.3.2. Accommodation Won't Help . . . 74

4.4. The Dynamic Strategy . . . 80

4.4.1. The Dynamics of Vagueness . . . 80

4.4.2. Dynamic Problems . . . 83

4.5. Other Alternatives . . . 86

4.5.1. The Diagonalization Strategy . . . 86

4.5.2. Standard Theories of Vagueness . . . 88

4.6. Semantic Quietism . . . 90

4.7. Quietism and Conversations . . . 92

4.7.1. When to Become Silent? . . . 95

4.7.2. Life Without Principle (II) . . . 97

Chapter 1

Introduction

My dissertation is the product of an obsession with something that doesn't exist: sharp cut-offs drawn by vague predicates. Let me explain. The characteristic mark of vague predicates—'tall', 'fast', 'rich'—is that they can be used to classify without drawing sharp boundaries. These classifications have positive cases on one side, negative cases on the other, and, yet, there is no last positive and no first negative case. This is what makes vague classifications cut-off free and what makes the phenomenon of vagueness remarkably puzzling. However, every substantial theory of vagueness on the market postulates sharp cut-offs left and right, up and down. As such, these theories idealize away the very phenomenon they seek to explain. I got obsessed with those cut-offs; all I wanted is a theory of vagueness free of them. Finding such a theory became the main target of my dissertation.

The upshot of my research is that any semantic theory that is compatible with the phenomenon of vagueness has to be, as a matter of principle, incomplete. There has to be some *robust* semantic facts that such a theory cannot capture. This is the only way in which a theory of vagueness can avoid cut-offs of an unacceptable kind. In Chapter 2 I motivate this view by arguing that any theory that aims to capture every robust semantic fact—in particular facts about the range of application of vague predicates—is forced to idealize vagueness away. In Chapter 3 I develop a semantic theory that is compatible with the phenomenon of vagueness; the theory is, of course, incomplete, but it is the only one that makes room for vagueness. The advantage of this theory is that it doesn't postulate sharp cut-offs separating different semantic categories. I call such a theory *Semantic Quietism.* Finally, in Chapter 4 I show how Quietism can be useful as part of a theory of linguistic communication. In particular I show how a broadly

Stalnakerian picture can incorporate Quietism's insights, thereby making room for the phenomenon of vague linguistic communication. In what follows I shall offer a brief description of each chapter.

Chapter 2: On Cutting Off Vagueness

First I asses theories that attempt to give a full account of the range of application of vague predicates. These theories typically offer a complete tripartite classification of the members of a sorites series for a vague predicate F, classifying members as positive, negative, or borderline cases.[1] The sense in which these theories aim for completeness is clear: they assign a semantic status to every member of a sorites series. This kind of approach is employed by, for example, Three-Valued Theory (Tappenden (1993),Soames (1999), and Tye (1994)), Supervaluatinism (Fine (1975), McGee & McLaughlin (1995), and Keefe (2000)), Degrees Theory (Machina (1976), Edgington (1997), and Smith (2008)), and Contextualism (Raffman (1994) and Shapiro (2003)). These theories face a well-known problem of higher-order vagueness. Vague predicates are such that there is no sharp cut-off between clear cut positive cases of their application and any other category. On the face of it, these theories postulate a sharp cut-off between positive and borderline cases, thereby raising Higher-Order Vagueness worries.

A leading strategy for trying to avoid Higher-Order Vagueness worries is to introduce a determinately operator, D. According to this approach, if Olivia is determinately tall, then it is (semantically) determinate that Olivia is tall, that is, it is the case that $DTall(Olivia)$. If she is determinately not tall, then it is determinate that she is not tall, that is, it is the case that $D\neg Tall(Olivia)$. And if she is a borderline case of a tall person, then it is indeterminate whether Olivia is tall; that is, it is the case that $\neg DTall(Olivia) \wedge \neg D\neg Tall(Olivia)$.

I offer an original proof to the effect that this strategy fails to avoid sharp cut-offs. The proof is stronger than other purported proofs in the literature since it appeals to more minimal assumptions. Other proofs of this kind assume that there are

[1] According to some theories this kind of classification comes in degrees, but this doesn't make a difference to the main argument.

some objects that are determinately borderline cases, or that some sentences are ω-determinately true (can be preceded by infinitely many determinately operators), or gap principles such as $\neg\exists x((D^n Bx \wedge D\neg D^{n-1}Bx'))$, or that $p \vDash Dp$ is valid. All these assumptions are controversial. The proof I offer doesn't appeal to any of them and also doesn't appeal to any assumptions that proponents of those proofs would reject. My proof is thus more compelling than the other would-be proofs, since it only assumes a KT modal logic, a Kripke-like semantics for modal operators, and that certain sentences of the form $\exists x(\neg D^n Fx \wedge \neg D\neg D^{n-1}Fx)$ (where D^n stands for n iterations of D) are D^m true for a finite m.

My proof shows that the popular theories of vagueness mentioned above are committed to the existence of sharp cut-offs. This is quite problematic, since there is no plausible explanation of what could possibly determine the exact location of the cut-off, *rather* than a slightly different one. This kind of consideration is what prompts us to think that vague predicates don't draw sharp cut-offs to begin with. For the same reason, this kind of consideration urges us to look with suspicion at any semantic theory that entails that there is a sharp cut-off. The moral I draw is that any attempt to offer a complete account of the range of application of vague predicates will misrepresent the phenomenon of vagueness by being committed to the existence of a sharp cut-off at some order. This is the way I defend the view according to which it is impossible for an adequate semantic theory of natural language to capture all semantic facts. There are natural language facts that any adequate semantic theory must be silent about.

Chapter 3: The Quietist's Gambilt

According to Quietism, the view that I defend, rather than offering a full account of the range of application of vague predicates we should offer an incomplete one; this is the best we can do. Thus, there are some objects such that our theories have to be completely silent as to whether they are positive, negative, or borderline cases of application of a vague predicate. Such silence, it's important to note, is compatible with every case being either a positive, negative or borderline case of application of the predicate. But, according to Quietism, it's a mistake for a semantic theory to try to cover every possible case. According to Quietism, we should cover enough cases to adequately satisfy our theoretical goals. For instance, if our goal is to model linguistic communication, Quietism urges us to cover enough cases to do so, without attempting to cover every single case.

Now, it's not the case that there is only one Quietistic model that can cover enough cases, relative to a fixed set of theoretical goals. The job of the semanticist is to identify one among many Quietistic models that are good enough for her purposes. There is no need to identify all of them—for higher-order vagueness reasons such a thing cannot be done. This makes the choice of a particular Quietist model a semi-arbitrary matter.

Thus, Quietism trades a cut-off between positive cases and borderline cases for one between positive cases and a silence range. A Quietist cut-off is, I argue, inoffensive. This is so because a Quietist would choose her silence region wisely. Doing so requires, among other things, becoming silent before covering *all* of the positive cases. If she does so, she guarantees that the cut-off between the positive and the silence range isn't between positive cases and cases that have some other semantic status. Instead the Quietist cut-off will land between two positive cases. This way of modeling the range of application of a vague predicate has the advantage of being compatible with the phenomenon of vagueness; no other view can do that. The disadvantage is that the model is silent about certain facts, and, as such, it cannot model the phenomenon completely. When compatibility with the phenomenon of vagueness is one of our main theoretical goals, we have to be Quietists. If we embark on a different project, one where vagueness isn't at issue, it is permissible to aim for completeness. One of the main themes in my project is that, unpleasant as it is, this is a trade-off that we are required to accept.

Quietism is in direct opposition to standard theories of vagueness, insofar as they aim for completeness.[2] Once this assumption is dropped, Quietism is compatible with the central components of the standard views. Crucially, Quietism is compatible with a modified version of Supervaluationism. I call this view 'Superquietism.' The objective

[2] Completeness is what I call 'exhaustiveness' in "The Quietist's Gambit".

of this theory is to have a supervaluationist-like theory that includes a region of silence between the determinate cases and the borderline cases. This can be done by modifying the notion of supertruth, the semantic clauses for the determinacy operator, and, crucially, the set of classical interpretations it quantifies over. The result is a system that is immune to Higher-Order Vagueness worries.

Superquietism preserves classical logic for the fragment of the language that is D-operator free. In order to get this result we need to assume a local notion of validity. It also preserves all penumbral connections for the same language fragment and it is able to offer essentially the same solution to the sorites paradox that supervaluationists offer. One important difference between Superquietism and Supervaluationism concerns operator D. Certain sentences containing D are valid in a Supervaluationist semantics, whereas Superquietism has to be silent about them. But that, I argue, is an acceptable price to pay. On a cost-benefit analysis, Superquietism should be preferred over Supervaluationism.

Chapter 4: Silence in Conversation

Quietism is a framework that many existent theories of language and communication can find useful. In my dissertation, I show how someone who wishes to offer a Stalnakerian model of linguistic communication could benefit from Quietistic insights. Insofar as number of the standing presuppositions in a given conversation are vague, a Stalnakerian approach to linguistic communication has to take into account the phenomenon of vagueness. This requires a non-trivial reassessment of how we think of the common ground, context set, and some pragmatic principles of rational communication.

As it stands, Stalnaker (1978) assumes the following principle of rational communication:

- (P) Any assertive utterance should express a proposition, relative to each possible world in the context set, and that proposition should have a truth-value in each possible world in the context set.

This principle is meant to restrict the effects of vagueness in discourse and to serve as a bridge principle between semantic and pragmatic presuppositions. If we accept (P), the story goes, there is no room for an explicit representation of vagueness in this theory of linguistic communication. Vague assertions, according to this view, trigger conversation crashes or pragmatic repairs, none of which require a representation of the context set that makes room for vagueness. This principle is widely assumed by Stalnakerians (Von Fintel (2008), Rayo (2008), Barker (2002)) either implicitly or explicitly. Von Fintel (2008) even takes (P) to capture an "irreducible property of natural language pragmatics." I argue that this Stalnakerian principle is based on shaky foundations. There is a battery of examples where felicitous vague assertions violate (P) without this having any major effect in the flow of the conversation. It seems, then, that we do update the context set with vague assertion, and that the result is a vague context set.

Several attempts to incorporate vagueness into the Stalnakerian framework are surveyed and found wanting. I conclude by showing how this framework can benefit from Quietism. In particular I show how the best way of making the Stalnakerian framework compatible with the phenomenon of vague linguistic communication is by being silent about certain features of the context set and the propositions expressed by vague assertions. To be more specific, the theory has to remain silent as to whether certain worlds are in the context set and whether certain worlds are in the propositions expressed by vague assertions. Finally I show how one can model the evolution of conversations given these limitations.

Chapter 2

On Cutting-Off Vagueness

2.1 An Adequacy Condition

There are many theories of vagueness out there. Most of them have been explored with all detail. None of them is free of serious objections. These objections have very different flavors. Some theorists criticize views that do not preserve classical logic. Others criticize views that do preserve classical logic. Some philosophers complain that according to some theories the major premise in the Sorites Paradox is true, while others complain that some other theories take it to be untrue. Some criticize theories that characterize vagueness as a semantic phenomenon, whereas others complain that certain theories don't do so. It is a shocking that there is such a wide disagreement about what constitutes an acceptable objection to a theory of vagueness. There seems to be no agreement about what exactly an *adequate* theory of vagueness must capture. It is as if the phenomenon hasn't even been clearly demarcated.

There is, however, one exception to this. Almost everyone seems to agree that an adequate theory of vagueness has to do justice to the fact that vague predicates do not draw sharp cut-offs between positive (or negative) instances of the application of that predicate and any other category. This comes as no surprise given the way we typically think about vagueness. Here is an example that illustrates this point:

Example 1: Good Runners

> You are observing the leading runners of today's 5k. They are very fast and in excellent shape. Of course, 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 are slow and out of shape. This is a nice sorites series. A friend approaches you and asks: 'did you have a chance to see good runners?' To which you reply: 'Yes, the fast ones are good runners.'

It is clear that based on your assertion you have used 'fast'—and 'good runner'—to draw a classification along the series of runners. You have classified some of them as fast and good runners, but not others. Now, it is also clear that based on your assertion we cannot find any good reason that would help us identify *the* last member of the series that has been classified as a good runner. The default position is that this is so because there is nothing like the last member of the series that is *determinately* a good runner.[1] The existence of such a member would entail that there is a pair such that the first member is determinately a good runner, but not the second, even though their running abilities are indistinguishable for all practical purposes. The existence of such a pair seems, *prima facie*, absurd.

Now, based on examples like Good Runners, several metaphors attempting to capture what vagueness is about suggest themselves. Frege's influential characterization of vague predicates is a good instance of this. In *Grundgesetze* ∮56 he explains the distinction between sharp and vague predicates in the following way:

> A definition of a concept (of a possible predicate) must be complete; it must unambiguously determine, as regards any object, whether or not it falls under the concept [...] We may express this metaphorically as follows: the concept must have a sharp boundary. If we represent concepts in extensions by areas on a plane, this is admittedly a picture that may be used only with caution, but here it can do us good service. To a concept without sharp boundary there would corresponds an area that had not a sharp boundary-line all round, but in places just vaguely faded away into the background.

This is quite metaphorical, but it is clear enough for our present purposes. The idea is that a sharp predicate—or concept—draws a precise line between the objects it applies

[1] Of course, Epistemicists are the notable exception.

to and the rest. For example, the predicate ‘is an even number’ draws a sharp cut-off between the even and odd numbers. A vague predicate, on the other hand, is very different in this respect. There is no sharp line dividing the instances of application from all the rest, but rather an area that fades away into the background. Any vague predicate—‘blond’, ‘near x’, ‘child’, ‘rich’, and so on—is a good example of this.[2]

Not everyone would feel comfortable with Frege’s metaphor. There are many metaphors out there, but at the end of the day all of them point in the same direction. Take, for example, the influential metaphor in Sainsbury (1996). According to him, vague concepts are concepts without boundaries. This captures right away the idea that we pursue; if there are no boundaries, there are no sharp cut-offs. Sainsbury offers a further metaphor to clarify, at least partially, what he has in mind:

> We could, perhaps, think of such concepts as like magnetic poles exerting various degrees of influence: some objects cluster firmly to one pole, some to another, and some, thought sensitive to the forces, join no cluster. (p.258)

It is not clear whether this metaphor is adequate, or whether it can be cashed out in a successful way.[3] What matters for now is the general spirit of the metaphor; it tries to capture the view that vague concepts do not draw sharp cut-offs.

Wright (1976) has also given voice to this kind of characterization of vagueness. His veiw is less metaphorical, and perhaps more theory laden, but it is quite useful to keep in mind nevertheless. On his view

> [...] no sharp distinction may be drawn between cases where it is definitely correct to apply such a predicate and cases of any other sort.

Here again is the view according to which vague predicates do not draw any sharp distinctions. This seems to be the core feature of vague predicates.

---

[2] I’m assuming, for the sake of simplicity, that vague predicates are sorites susceptible. It has been argued (Weatherson (2010)) that not all vague predicates are like this. If this is correct, then this paper is only concerned with a subclass of vague predicates—those that are sorites susceptible.

[3] The reason for doubt is that this metaphor draws a sharp distinction between the objects that are firmly attached to one pole and the ones that aren’t. Vague predicates don’t draw sharp distinctions like that.

These ways of thinking about vagueness are motivated by the observation that a very small difference along the relevant dimension cannot affect the justice with which vague predicates are applied. In Wright's words, we say that a vague predicate is *tolerant* when it satisfy this condition.[4] For example, the predicate 'tall' is tolerant because a small difference in height doesn't affect the justice with which the predicate is applied—one millimeter cannot make the difference between someone who is tall and someone who isn't. It is this kind of observation that motivates the idea that vague predicates don't draw sharp distinctions, boundaries, or cut-offs. Thus, without attempting to define the concept of vagueness, we can be confident that vague predicates are used to draw *vague classifications* that meet the following two conditions:

(a) Relative to a suitable sorites series vague classification have some positive and some negative cases.

(b) There is no last positive case and no first negative case—assuming that the series is arranged from positive to negative.[5]

The second feature is motivated by the observation that vague predicates don't draw sharp cut-offs. This is what makes the phenomenon of vagueness extremely puzzling. If there are positive cases on one end and negative case on the other, how can it be that there is no last positive case?[6] If there is no last positive case, how can there be a single negative case? These kinds of considerations are the ones that fuel to the Sorites Paradox.

At the very least, an adequate theory of vagueness has to be compatible with the fact that vague predicates do not draw sharp cut-offs. A theory of vagueness that isn't compatible with this simply doesn't make room for the phenomenon it's trying to account for. Given this, a *condition of adequacy* for a theory of vagueness is that its representation of the phenomenon does justice to the fact that vague predicates don't

[4]This way of thinking about tolerance can be found in Wright (1976). For similar notions of tolerance see Shapiro (2006) and Rayo (2008)

[5]Depending on your theoretical inclinations you may one to interpret (b) as saying that there is no last *determinate* positive and no first *determinate* negative case.

[6]As before, you can read this as concerning determinate positive and negative cases, if you like.

draw sharp cut-offs of any kind. This isn't the only condition of adequacy, but it is a basic one.

Let me mention an important caveat before we move on. There can be theories that capture certain features of vague sentences without satisfying our basic condition of adequacy. Let's assume, for the sake of illustrating this point, that vague predicates have borderline cases. As such, some vague sentences can be neither true nor false. A Three-Valued Theory can do an excellent job at capturing this purported fact. However, this isn't to say that just because Three-Valued Theory can do this, it is an adequate theory of vagueness. This theory captures one aspect of the phenomenon, but it fails to capture central aspects of vagueness that distinguish it from other phenomena. In particular, it fails to capture the fact that vague predicates don't draw sharp cut-offs.[7] An central theme in this paper is that many theories of vagueness adequately capture many aspects of the phenomenon of vagueness, but that they fail to capture one essential feature—the lack of cut-offs.[8]

In what follows I shall argue that all the known theories of vagueness that are in the business of modeling the range of application of vague predicates fail to meet this basic condition of adequacy. They all represent vague predicates as drawing some sort of sharp cut-off. By doing so they cut-off vagueness altogether. My diagnosis is that they make this mistake because they adopt a very specific semantic assumption. I call this assumption *Exhaustiveness*.

2.2 Exhaustive Theories of Vagueness

At a certain level of abstraction, theories of vagueness that are in the business of modeling the range of application of vague predicates fall into one of two categories. Either they model the range of application in terms of positive, borderline, and negative cases, or they do so *only* in terms of positive and negative cases. If the predicate applies to an object, then that object is a positive case, if it doesn't apply, then it

[7]Of course, I can give substance to this claim after discussing Higher-Order Vagueness.

[8]Thanks to Van McGee (p.c.) for a discussion that lead to these remarks.

is a negative case, and if this is—to some degree—indeterminate, then we say that it is a borderline case. Theories in the first category take vagueness to be a semantic phenomenon. Notable theories of this kind are Supervaluationism (Fine (1975), Kamp (1975), McGee & McLaughlin (1995), and Keefe (2000)), Degree Theory (Machina (1976), Edgington (1997), Weatherson (2005), Smith (2008)), Three-Valued Theory (Tye (1994), Tappenden (1993), Soames (1999), and Contextualism ((Shapiro, 2003), Kamp (1981b), Fara Graff (2000) Barker (2002), Raffman (2014)). Theories in the second category take it to be an epistemic phenomenon. The notable theory of this kind is Epistemicism (Williamson (1999), Sorensen (2001), Kennedy (2007)). Throughout this paper I shall only focus on theories of the second kind. For excellent criticisms of theories of the first kind see Gómez-Torrente (2002) and Fara Graff (2002). Now, this way of grouping theories isn't the most perspicuous one, given that some of this theories are different enough from other theories in their group to deserve a group of their own—two glaring examples are Fara Graff (2000) and Weatherson (2005). I've grouped theories this way for the sake of simplicity. As we shall see what matters for our current discussion isn't that this-or-that theory is in such-and-such a group. What matters is that all of them, and many others, accept Exhaustiveness.[9]

To a first approximation, Exhaustiveness is the semantic assumption requiring that each member of the domain must have a full semantic profile. This assumption, I take it, is motivated by a desire to capture every robust semantic fact and to achieve maximum logical rigor—without this assumption it isn't clear that we can satisfy both desires. As we shall see, the phenomenon of vagueness isn't compatible with Exhaustiveness. Consequently, any theory that adopts this assumption won't be compatible with the

[9]There are at least other two other kinds of theory. Some theorists, let's call the the *gapists* hold that vague predicates fail to have a range of application relative to soritical domains—hey only have a range of application relative to domains that have a significantly large gap between positive and negative cases. This kind of theory is defended by Manor (2006), Rayo (2008), Gómez-Torrente (2010), and Pagin (2010). The reason why I won't discuss this kind of theory any further is because it denies the phenomenon that this paper focuses on—the absence of cut-offs drawn by vague predicates relative to a soritical domain. The other kind of theory, call it Homophonic Theory, isn't even in the business of modeling the range of application of predicates. At most what this kind of theory says is things like: 'red' is true of something iff that thing is red. This kind of theory is represented by Sainsbury (1996) and Ludlow (2014). Given that this theory isn't in the business of modeling the range of application of predicates, I won't discuss it any further.

phenomenon of vagueness.

Here is a more precise formulation of this semantic assumption:

- **Exhaustiveness**: For each object o in the domain and for each predicate $\ulcorner\Phi\urcorner$ in the language, the semantic status of o with respect to $\ulcorner\Phi\urcorner$ is fully specified; either $\ulcorner\Phi\urcorner$ applies to o, or it doesn't, or this is, to some degree, indeterminate.

Every semantic model on the market adopts this assumption; this is so whether the predicate in question is $\ulcorner Tall(x)\urcorner$, $\ulcorner DTall(x)\urcorner$, $\ulcorner D\neg DTall(x)\urcorner$, $\ulcorner\neg DD\neg DTall(x)\urcorner$, or what have you. These theories will always tell you whether this predicate applies to the object or not, or whether this is indeterminate. At face value, this looks like a highly reasonable assumption; if one wants to offer a complete and rigorous description of a semantic or epistemic phenomenon one's theory needs to take a stand on every single case. So why make an exception when theorizing about the phenomenon of vagueness? If we want to capture every singe robust aspect of the phenomenon, Exhaustiveness may appear to be a non-negotiable assumption. As we shall see, the root of the problem is with this theoretical ambition.

2.3 Cut-Offs, Cut-Offs Everywhere

When vagueness isn't at issue Exhaustiveness feels extremely natural, but when our theoretical goals shift towards the phenomenon of vagueness, this assumptions looks very suspicious. The reason for this is simple; an exhaustive semantic model is maximally precise. If you ask a question to an exhaustive model about the semantic status of a given object relative to a predicate, the model will deliver a sharp answer—the model will predict that the predicate applies, or that it doesn't, or that this is indeterminate, or that it is indeterminate whether it is determinate, or what have you. Nevertheless, one may think that if the model is in fact making room for the phenomenon of vagueness, the model itself cannot be so sharp.

A concrete way of seen what's wrong with exhaustive models is to consider how they represent the range of application of vague predicates relative to a suitable sorites series. If the model is exhaustive, then there has to be a pair of adjacent objects with different semantic profiles; there has to be a sharp transition from the cases where the predicate applies to those where there is a different semantic status. Otherwise the model could only have positive cases of application of the predicate. This sharp semantic transition is a cut-off, and vague predicates do not draw them. There is no way around it—that's is the nature of the semantic tools that we use.

Of course, things aren't quite this simple. An Exhaustiveness defender may argue that the kind of criticism overlooks the phenomenon of Higher-Order Vagueness. This line of defense typically uses a determinacy operator (D) to argue that, in these exhaustive models, there is no sharp division between the clear cases and the borderline cases, and the negative cases and the borderline cases. Here is how the argument goes. If $DTall(Maria)$ then we say that Maria is determinately tall. If $D\neg Tall(Maria)$ we say that Maria is determinately not tall. And if $\neg DTall(Maria)$ and $\neg D\neg Tall(Maria)$ we say that Maria is borderline tall. Exhaustiveness forces us to say that there is a pair of adjacent members of the series, let's say Olivia and Hugo, such that $DTall(Olivia)$ and $\neg DTall(Hugo) \wedge \neg D\neg Tall(Hugo)$. This is so even if the difference in height between Olivia and Hugo is incredibly small. This looks like a sharp cut-off, but this is something that we already noted.

However, the argument continues, there is no sharp cut-off between the determinate cases and the borderline cases (between Olivia and Hugo) because even though Olivia satisfies $DTall(x)$, she also satisfies $\neg DDTall(x) \wedge \neg D\neg DTall(x)$. Similarly, even if Hugo satisfies $\neg DTall(x)$, he also satisfies $\neg DDTall(x) \wedge \neg D\neg DTall(x)$. Thus, the argument goes, the cut-off isn't determinately there; neither Olivia is determinately determinately tall, nor Hugo is determinately borderline tall. They are second-order borderline cases.[10] The strategy is, then, to move up one order, look down, and see that the cut-off isn't determinately there. The idea is to keep doing this until we are satisfied.

It doesn't take a very sophisticated argument to be suspicious about these moves. The first observation has been adequately captured by Sainsbury's colorful words: 'you

[10] For solid arguments against the coherence of the notion of higher-order borderline case see Wright (2003) and Raffman (2014).

do not improve a bad idea by iterating it'. The idea, I take it, is that if it is problematic to model vagueness in terms of positive, borderline, and negative cases, applying the same idea at higher orders won't make the theory less problematic. Before these theories had the problem of justifying why according to their models $DTall(Olivia)$ and $\neg DTall(Hugo) \wedge \neg D\neg Tall(Hugo)$ even though the difference in height between them is insignificant. Now these theories also have the problem of explaining why some other pair of adjacent members in the series, say $\langle Andrea, Alejandra\rangle$, is such that $\ulcorner DDTall(Andrea)\urcorner$ and, even if the difference in height between them is insignificant, $\ulcorner\neg DDTall(Alejandra)\urcorner$. A similar question will come up for each time this kind of theory gets iterated.

Now, consider the final super-structure of orders of vagueness that a theory like this would offer. If you could see it, it would look like a huge skyscraper where each floor contains cut-offs trying to eliminate the cut-offs in the floor immediately bellow. Each of these cut-offs land in a particular location along the relevant sorites series—either between Oliva and Hugo, Andrea and Alejandra, or some other pair. The question now is, why this particular skyscraper rather than a slightly different one? How can it be that our linguistic practices in conjunction with the way the world is justifies this particular super-structure rather than one where all the cut-offs are moved one place to the right? One central reason why we reject the existence of a sharp cut-off between positive cases and any other category is that noting in our linguistic practices and the way the world is could justify it. But then, absolutely nothing in our linguistic practice and the way the world is could justify one of this maximally precise super-structures rather than a slightly different one. It seems, then, that an appeal to higher-order vagueness as an attempt to save exhaustive models is misguided.

2.4 Higher Problems

We have presented so far a very intuitive argument against exhaustive models of vagueness. Now it is time to put on the table a technical argument. The point is a familiar one. A finite sorites series can only support a finite number of distinctions, so long as for each distinction there is a member of the series that definitely falls under it.

Thus, assume that the first member of the series is $D^{\omega}B(0)$, and the last member is $D^{\omega}\neg B(100)$—where $\ulcorner B \urcorner$ is vague, D^{ω} stands for infinite iterations of D, and each member of the series gets assigned a number between 0 and 100 depending on where they sand in the series.[11] Now, given that $\ulcorner B \urcorner$ is vague, this kind of theory tells us that there are first-order borderline cases. Let's call one of these borderline cases $\ulcorner n \urcorner$. Now, let's say that n is determinately a first-order borderline case of the predicate $\ulcorner B \urcorner$. Thus, $D(\neg DB(n) \wedge \neg D\neg B(n))$. To make room for second order vagueness, we are told that there are second-order borderline cases. Lets call one of them $\ulcorner m \urcorner$, and let's assume that m is determinately second-order borderline B. Thus, we have that $D(\neg DDB(m) \wedge \neg D\neg DB(m))$. We know that $m \leq n$, given that n is determinately a first-order borderline case, and m is determinately a second-order borderline case towards the positive side of the series (recall that the numbers closer to 0 are positive and the ones closer to 100 are negative cases).[12] Now, clearly $n \neq m$, given that $D\neg DB(n)$ and $\neg D\neg DB(m)$. Therefore, $m < n$. A similar reasoning can be used to show that there is a member of the series, call it $\ulcorner o \urcorner$, that is determinately a third-order borderline case, and that $o < m$. Now, this kind of theory needs to keep postulating borderline cases at every order to avoid sharp cut-offs—this is how they treat the phenomenon of vagueness. However, there is a limit as to how high they can go. The limit is imposed by 0, who cannot be a borderline case of $\ulcorner B \urcorner$ at any order, given that $D^{\omega}B(0)$. For simplicity, let's assume that 0 is the only member of the series that is $D^{\omega}B(0)$. Thus, 1 is a determinate borderline case of $\ulcorner B \urcorner$ at some really high order. Thus, for some n, $D(\neg DD^{n}B(1) \wedge \neg D\neg D^{n}B(1))$. From this it follows that $D\neg DDnB(1)$. But we know that $DDD^{n}B(0)$. Thus, $D(DD^{n}B(0) \wedge \neg DD^{n}B(1))$. This is a sharp cut-off that directly clashes with the nature of vagueness. After all, it is natural to think that nothing could possibly justify the exact location of that cut-off rather than a slightly different one.

This kind of theory postulates higher-order borderline cases to avoid the presence

[11] Thus, for example, if $D^{\omega}B(0)$ we say that 0 is absolutely bald, or whatever the interpretation of $\ulcorner B \urcorner$ is. And we say that the name of 0 is $\ulcorner 0 \urcorner$', the name of 1 is $\ulcorner 1 \urcorner$, and so on.

[12] A second-order borderline case towards the negative side of the series satisfies $D(\neg DD\neg B(x) \wedge \neg D\neg D\neg B(x))$.

of sharp cut-offs. We have seen that, given certain assumptions, there is at least one sharp cut-off that is unavoidable.[13] Now, if it is unbelievable that our conventions and linguistic practices determine a cut-off at the first order level, it is even harder to believe that they determine a cut-off at such a high-order. It is hard to believe that our linguistic practices determine anything at all at really high orders. What could possibly determine that the cut off is there, rather than at some nearby location? If no plausible answer can be given to this question, it should be easier to give a more plausible argument to the effect that there is a first-order cut-off, but this would render the whole view untenable.

Our presentation of this argument assumed that there are definite borderline cases at every order. Given this assumption we could easily prove that each order requires a definite borderline case that is different from every other definite borderline case at other orders. Thus, the idea was that there has to be a definite first-order borderline case, and a definite second-order borderline case, and a definite third-order borderline case, and so on, and that each of them is different from each other. If this is the case, then it is easy to show that a finite series can only hold a finite number of distinctions.

It could be objected to this kind of argument by pointing out that vagueness doesn't require that there are definite borderline cases, it only requires that it is definite that there are borderline cases.[14] Thus, instead of demanding that $\exists x D(\neg DB(x) \wedge \neg D\neg B(x))$, this view only requires that $D\exists x(\neg DB(x) \wedge \neg D\neg B(x))$, and the same goes for second and higher orders of vagueness. In the argument above it was crucial to distribute the determinacy operator over the conjunction, but if the operator has wide scope over the quantifier, this crucial move has been blocked. Thus, our argument hasn't shown that there has to be a cut-off at some order.

This kind of move has an independent motivation. This is how the argument goes: It is clear that some predicates have determinate borderline cases. Take the case of the artificially defined predicate 'domal'. We can stipulate that every domal is a mammal

[13] We have shown that much, but, of course, with the same set of assumptions we could show that there are many more.

[14] This objection was suggested to me by Brian McLaughlin (p.c).

and that every dog is a domal. Nothing else is stipulated regarding this predicate. Now, it is clear that Bisha, the cat, is determinately a borderline case of 'domal'. Notice that this predicate isn't vague. Now, legitimate vague predicates are such that it isn't that easy to identify its borderline cases. We could point at someone who is 1.70m tall and say that she is borderline tall. But, is it clear that she is determinately a borderline case? The answer to this question isn't clear. There seems to be no theoretical cost in saying that she isn't a determinate borderline case, and that in fact there are no determinate borderline cases at all. The same thing can be said for every borderline case at every order. Given that it is absolutely required that there are no sharp cut-offs at any order, we have to say that it is determinate that there are borderline cases at every order. But this doesn't mean that there are objects that are determinate borderline cases.

If this view is correct, then our argument has effectively been blocked. However, it isn't clear that even this view is compatible with the absence of cut-offs at some order. Initially one may think that it is. In fact, by paying attention to the following simple model one could think that in fact there can be borderline cases at every order relative to a finite Sorites series:

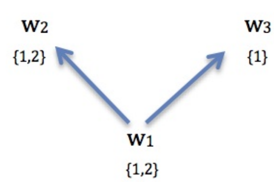

Let the sets bellow each w represent the extension of $\ulcorner B \urcorner$ at that world, and assume that the only members in the domain are $\{1, 2, 3\}$. Also, assume that reflexivity holds. For simplicity, let's call a sentence of the form $\exists x(\ DD^n B(x) \wedge\ D\ D^n B(x))$, a borderline sentence. Now, given that in w_2 it is true that $D^\omega B(2)$ and at w_3 it is true that $D^\omega \neg B(2)$, in w_1 it is true that for every n, $\exists x(\neg DD^n B(x) \wedge \neg D \neg D^n B(x))$—of course, it is 2 the one satisfying this existential at w_1 for every n. What this toy model shows is that there is a finite series of objects that supports the truth of an infinite number of borderline sentences at w_1. This seems to be an argument to the effect that a finite series can hold an infinite number of distinctions, and that, therefore, we don't have to be committed to a cut-off at any order. At this point it may be tempting to think that the reason why we thought that vagueness has to run out at some order is because we mistakenly thought that vagueness requires the existence of determinate borderline cases. If we jettison this requirement, one may think, we can successfully show that there is no sharp cut-off at any order.

However, this kind of argument fails. The first thing to note is that in this toy model no borderline sentence is definitely true; this is so given that there are no borderline cases according to w_2 and w_3. But surely it is determinately true that there are borderline cases. Furthermore, there doesn't seem to be any good reason to doubt that it is D^ω true that there are borderline cases; it is natural to think that it is absolutely settled that there are no cut-offs, if language is in fact vague. Thus, this model fails to show anything interesting about borderlines. The only thing that this toy model shows is that there is a tremendous mismatch between actual vagueness and the truth of plain borderline sentences.

It shouldn't come as a surprise that this toy model isn't completely successful; it is too simple to capture such a tremendously complex phenomenon. However, is there an exhaustive model with only a finite number of objects in its domain that verifies all the borderline sentences along with their D^ω versions? In the next section I shall lay out an argument demonstrating that there isn't such a model. The next section is devoted to showing precisely that.

2.5 Vagueness Must Stop

In this section I shall argue that, given some minimal assumptions, exhaustive models are committed to the existence of sharp cut-offs. In particular I shall offer a proof to the effect that given these minimal assumptions there is a sharp cut-off at some order. To be even more precise, I proof that for some n, and for some vague predicate $\ulcorner F \urcorner$, the following sentence is true: $\neg\exists(\neg D^n Fx \wedge \neg D\neg D^{n-1}Fx)$. The final proof is stronger than other proofs (Wright (1987), Fara (2003), Fine (2008), Zardini (2013)) because it doesn't need to assume that some sentences Φ are $D^\omega\Phi$, nor it needs to assume that $p \vDash Dp$, $Dp \to DDp$, or any gap principle $(\neg\exists x((D^n Bx \wedge D\neg D^{n-1}Bx')))$.[15] If the proof that I'm about to present has a weakness, it is that forces theorists of Higher-Order Vagueness theorists into a particular modal theory of reference—it forces them to interpret D in terms of Kripke modal models.[16] Given that Kripke models (Kripke (1963)) are by far the best and most popular way of interpreting this kind of operator, this is a minor price to play.

The argument relies on the following assumptions:

1. The accessibility relation ('R') is reflexive.

2. We employ a constant domain semantics (thus, if an object exists at a world in the model, then that object exists in every other world in that model.)

3. Leibniz's Law

4. The semantic rule according to which if a sentence of the form $\ulcorner\exists x\Phi x\urcorner$ is true at w, then there is an object o in w that satisfies $\ulcorner\Phi x\urcorner$.

The first assumption is not negotiable, and I don't see any point in rejecting the second one. Notice that assumption (2) is reasonable given that what we are trying to model is the phenomenon of vagueness relative to a *given* sorites series. As such, the existence of the members of the series should't vary across possible worlds; what should vary are the extensions of the relevant predicates. Now, given (1) we only need to assume that the logic of D is KT. It would take a paraconsistent logic to deny (3) and an unreasonable semantics for the existential quantifier to deny (4).

For the sake of simplicity let's assume that the only elements in our domain are the members of a *finite* sorities series for vague $\ulcorner F\urcorner$. Also, in what follows I will use

[15]It's worth pointing out that the other proofs I mention attempt to show that Higher-Order Vagueness leads to contradiction. However, these proofs could also be used to show that there are sharp cut-offs at high-orders. Thanks to Brian Weatherson for pointing this out.

[16]Thanks to Vann McGee for pointing this out.

$\Box$ rather than D. The only reason I do this is because it is easier for me to follow the proof this way. Needless to say, $\Diamond\Phi$ gets defined as $\neg\Box\neg\Phi$. Finally, let's call a *borderline sentence* a sentence of the form $\exists(\neg\Box^n Fx \wedge \neg\Box\neg\Box^{n-1}Fx)$. Of course, there are borderline sentences with many different forms, but this is the particular form we shall focus on.

The first version of this proof assumes that if a borderline sentence Φ is true, then $\Box^\omega\Phi$—omega stands for an infinite iteration of boxes. The second version of this proof doesn't depend on this assumption—this is shown in Proof 3.

Consider the following borderline sentences:

- Bor_1: $\Box^\omega\exists x(\neg\Box Fx \wedge \Diamond Fx)$
- Bor_2: $\Box^\omega\exists x(\neg\Box\Box Fx \wedge \Diamond\Box Fx)$
- Bor_3: $\Box^\omega\exists x(\neg\Box\Box\Box Fx \wedge \Diamond\Box\Box Bx)$
- .
- .
- .
- Bor_{n-2}: $\Box^\omega\exists x(\neg\Box^{n-2}Fx \wedge \Diamond\Box^{n-3}Fx)$
- Bor_{n-1}: $\Box^\omega\exists x(\neg\Box^{n-1}Fx \wedge \Diamond\Box^{n-2}Fx)$
- Bor_n: $\Box^\omega\exists x(\neg\Box^{n}Fx \wedge \Diamond\Box^{n-1}Fx)$

Notice that '*Bor*' stands for 'borderline' and the subscript corresponds to the order. Thus, for example, Bor_3 states that it is omega determinate that there is a borderline case at the third-order.

Given this, we can prove the following (recall that the domain only contains members of a finite sorites series):

1. For every n, if $w \vDash Bor_1, ..., Bor_n$ then w's domain must contain at least n objects.

In order to get the flavor of how this proof works, let's first show that if $w \vDash Bor_1, ..., Bor_3$ then there are at least three objects in w's domain.

> *Proof 1*: Assume that $w \vDash Bor_1 \wedge Bor_2 \wedge Bor_3$. Now, given that $w \vDash Bor_3$, there has to be an object $o_1 \in w$ such that $w \vDash \Diamond\Box\Box Fo_1$. Thus, there is a w', such that wRw' and $w' \vDash \Box\Box Fo_1$. Given that wRw' we have that $w' \vDash Bor_2$.[17] Thus, there has to be an object o_2 such that $w' \vDash \neg\Box\Box Fo_2 \wedge \Diamond\Box Fo_2$. Therefore, $w' \vDash \Box\Box Fo_1 \wedge \neg\Box\Box Fo_2$. By Leibniz Law $o_1 \neq o_2$, so there are at least two objects in w. Now, given that $w' \vDash \Diamond\Box Fo_2$, there is a w'' such that $w'Rw''$ and $w'' \vDash \Box Fo_2 \wedge \Box Fo_1$. Since $w' \vDash Bor_1$ and $w'Rw''$, we know that $w'' \vDash \exists x(\neg\Box Fx \wedge \Diamond Fx)$. Then, there has to be an object $o_3 \in w''$, such that $\neg\Box Fo_3$. By Leibniz Law, $o_3 \neq o_1$ and $o_3 \neq o_2$, and given that $o_1 \neq o_2$, we have shown that there are at least three objects in w.

Now let's show that for ever n, if $w \vDash Bor_1, ..., Bor_n$, then there are at least n objects in w—notice how Proof 1 is very similar in important respects.

> *Proof 2*: Assume that $w \vDash Bor_1, ..., Bor_n$. Now, given that $w \vDash Bor_n$, there has to be an object $o_1 \in w$ such that $w \vDash \Diamond\Box^{n-1} Fo_1$. Therefore, there is a w', such that wRw' and $w' \vDash \Box^{n-1} Fo_1$. Given that wRw', we have that $w' \vDash Bor_{n-1}$. So there is an object o_2 in w' such that $w' \vDash \neg\Box^{n-1} Fo_2$. By Leibniz Law $o_1 \neq o_2$. Thus, given that the domain is constant, there are at least two objects in w. Now, we also have that $w' \vDash \Diamond\Box^{n-2} Fo_2$, so there is a w'', such that $w'Rw''$, and $w'' \vDash \Box^{n-2} Fo_1 \wedge \Box^{n-2} Fo_2$. But $w'' \vDash Bor_{n-2}$, so there is an object o_3 such that $w'' \vDash \neg\Box^{n-2} Fo_3$. By Leibniz Law $o_3 \neq o_1$ and $o_3 \neq o_2$, so there are at least three objects in w. Repeated application of this process shows that there is a $w*$, such that $w* \vDash Bor_1$. But we also know that $w* \vDash \Box Fo_1 \wedge \Box Fo_2 \wedge \Box Fo_3, ..., \Box Fo_{n-1}$.[18] Given this, there is an object

[17] Notice that if $\Box^\omega \Phi$ is true at w, then it is also true in every world that bares the ancestral of the accessibility relation to w.

[18] We know this because $w*$ is connected through the ancestral of the accessibility relation to a chain starting at w.

o_n such that $w* \vDash \neg\Box F o_n$. By Leibniz Law, $o_n \neq o_1 \wedge o_n \neq o_2, ..., o_n \neq o_{n-1}$. Therefore there are at least n objects in w.

We can use this result to show that if there is a sorites series with n objects, Bor_m (for $m > n$) has to be false in w, provided that $Bor_1, ..., Bor_n$ are all true in w. This result is already very problematic. If there is in fact vagueness at every order, what could possibly justify the failure of all sentences of the form Bor_m.

Someone complain that the assumption that some borderline sentences are $\Box^\omega$ is too strong. Dorr (2010), for example, has argued that no sentence $\ulcorner\Phi\urcorner$ is $\Box^\omega\Phi$. Independently of how plausible these arguments are, we can show that there is a weakening of our previous proof that doesn't require any borderline sentences to be $\Box^\omega$. This weakening shows that some unboxed borderline sentences are false, relative to a finite Sorites series. Thus, this kind of worry should be dismissed.

As before, we assume for simplicity that our domain only contains the elements of a finite sorites series for vague predicate F. Now, consider the following borderline sentences—where Bor_m differs from BOR_m only in the number of boxes in front of the existential quantifier.

- BOR_1: $\Box^{n-1}\exists x(\neg\Box Fx \wedge \Diamond Fx)$
- BOR_2: $\Box^{n-2}\exists x(\neg\Box\Box Fx \wedge \Diamond\Box Fx)$
- BOR_3: $\Box^{n-3}\exists x(\neg\Box\Box\Box Fx \wedge \Diamond\Box\Box Fx)$
- .
- .
- .
- BOR_{n-2}: $\Box\Box\exists x(\neg\Box^{n-2}Fx \wedge \Diamond\Box^{n-3}Fx)$
- BOR_{n-1}: $\Box\exists x(\neg\Box^{n-1}Fx \wedge \Diamond\Box^{n-2}Fx)$
- BOR_n: $\exists x(\neg\Box^{n}Fx \wedge \Diamond\Box^{n-1}Fx)$

The difference between these group of borderline sentences (BOR) and the one before (Bor) is that as we go up the orders less and less boxes are iterated—we run out of boxes at the nth order. Given this, we can prove the following weakening of (1):

2. For every n, if $w \vDash BOR_1, ..., BOR_n$, then w's domain contains at least n objects.

The following definition shall be of use throughout the proof:

- *Definition*: w' is n steps away from w just in case there is a chain containing $n+1$ worlds connected by the accessibility relation starting at w and ending at w'. Thus, for example, if wRw' and $w'Rw''$, we say that w' is one step away from w and w'' is two steps away from w.

Now we can proceed to prove (2):

Proof 3: Assume that $w \vDash BOR_1, ..., BOR_n$. Now, given that $w \vDash BOR_n$, there has to be an object $o_1 \in w$ such that $w \vDash \Diamond\Box^{n-1}Fo_1$. Therefore, there is a w', such that wRw' and $w' \vDash \Box^{n-1}Fo_1$. But we know that $w \vDash BOR_{n-1}$, so given that w' is at most one step away from w, $w' \vDash \exists x(\neg\Box^{n-1}Fx \wedge \Diamond\Box^{n-2}Fx)$. Thus, there is an object o_2 in w' such that $w \vDash \neg\Box^{n-1}Fo_2$. By Leibniz Law, $o_1 \neq o_2$. Given that the domain is constant, there are at least two objects in w. Now, we also have that $w' \vDash \Diamond\Box^{n-2}Fo_2$, so there is a w'', such that $w'Rw''$, and $w'' \vDash \Box^{n-2}Fo_2 \wedge \Box^{n-2}Fo_1$.[19] But we also know that $w \vDash BOR_{n-2}$, so, given that w'' is at most two steps away from w, $w'' \vdash \exists x(\neg\Box^{n-2}Fx \wedge \Diamond\Box^{n-3}Fx)$. So there is an object o_3 such that $w'' \vDash \neg\Box^{n-2}Fo_3$. By Leibniz Law, $o_3 \neq o_1$ and $o_3 \neq o_2$. So there are at least three objects in w. Repeated application of this process shows that there is a $w*$ that is at most n steps away from w such that $w* \vDash \exists x(\neg\Box Fx \wedge \Diamond Fx)$.[20] Given this, there is an object o_n such that $w* \vDash \neg\Box Fo_n$. But it is also the

[19]We get that $w'' \vDash \Box^{n-2}Fo_1$, because $w' \vDash \Box^{n-1}Fo_1$ and $w'Rw''$.

[20]This is so because $w \vDash BOR_1$ and $w*$ is at most n steps away from w, so we can get from w to $w*$ by discharging all the boxes (which are exactly n boxes) in BOR_1.

case that $w* \vDash \Box Fo_1 \wedge \Box Fo_2 \wedge Fo_3 \wedge, ..., \wedge \Box Fo_{n-1}$.[21] So, clearly, $o_n \neq o_1$, $o_n \neq o_2$, $o_n \neq o_3$,..., and $o_n \neq o_{n-1}$. Therefore there are at least n objects in w.

From this result it follows that if there are only $n-1$ objects in the domain—and, thus, in the relevant sorites series—and $w \vDash BOR_1, ..., BOR_{n-1}$, then BOR_n is false at w. Also notice that, given that the domain is constant, BOR_n is false at every world where all of $BOR_1, ..., BOR_{n-1}$ are true. Given that it is hard to believe that any of this sentences are false at a world in the model—after all, what could possibly falsify them?—it is hard to believe that BOR_n is something other than *determinately* false. There is no way around it, exhaustive models are committed to unjustifiable cut-offs.

2.6 Instruments, Artifacts, and Tools

We have considered exhaustive theories of vagueness. All of them are plausible to some extent, but it is also the case that all of them suffer from an important problem; they are all committed to the existence of cut-offs. This problem is a serious one because any theory of vagueness that has this commitment introduces precision when there is supposed to be vagueness, thereby idealizing away the phenomenon they are supposed to model. Adopting exhaustiveness is an effective way to cut off vagueness. As such, exhaustive theories of vagueness fail to meet a basic condition of adequacy. A correct diagnostic is that these theories of vagueness cannot aim for completeness, if they attempt to be, at the very least, compatible with the phenomenon of vagueness. The exhaustive is well suited to model what is precise, but not what is vague.

There is only one line of defense for those who think that there is nothing in principle wrong with using exhaustive mathematical resources to model a vague language. One must think that that the precision that comes with those mathematical resources doesn't really eliminate the phenomenon one is trying to model. To the best of my knowledge, the only reasonable way in which one can support this view is by arguing that the

[21]This is so because $w' \vDash \Box^{n-1} Fo_1$ and w' is at most n-1 steps away from $w*$, and $w'' \vDash \Box^{n-2} Fo_2$ and w'' is at most n-2 steps away from $w*$, and $w''' \vDash \Box^{n-3} Fo_3$ and w''' is at most n-3 steps away from $w*$, and so on.

proper way of thinking about the usage of those mathematical resources is as some kind of useful idealization (notable proponents of this view are Edgington (1997), Cook (2002), and Rayo (2008)). If this is correct, then perhaps we shouldn't worry about those cut-offs; they are just side effects of useful idealizations, and it is a mistake to take them seriously. This line of defense can take several forms. One can think about these mathematical resources as idealizations, tools, instruments, or artifacts. However one describes them, this view holds that a proper use of those idealizations doesn't get in the way of theorizing about vagueness.

As it turns out, how to articulate this line of defense isn't a trivial matter. On the one hand this view contains a very plausible component: it is desirable to introduce some idealizations in our models, or to think of certain aspects of our theories instrumentally. On the other hand, this is not to say that any aspect of our theories can be treated in this way; this would be just a license to do whatever we want. If one is to idealize, one must do it responsibly. What we need is a clear idea of which aspects of our theories we can take seriously and which ones we shouldn't. Once this is clear, it is a further question whether this distinction can be used to solve Higher-Order Vagueness worries. In what follows I shall consider an attempt to flesh out with some detail this particular kind of view.

Cook (2002) has the only attempt to justify with some detail the view according to which the precision that comes with standard mathematical resources shouldn't be taken seriously when modeling vague languages. On his view, a theory of vagueness shouldn't be thought of as a realistic description of the relevant phenomena—that would be an unrealistic request. Rather, the proper way of thinking about theories of vagueness is as models. The first thing to note about models is that they, unlike realistic descriptions, are intended as 'merely one tool among many that can further our understanding of the discourse in question.'(Cook (2002), p.234) As such '[I]n building models it is often advantageous (and sometimes unavoidable) to introduce some simplification. The idea is that we can eliminate, or at least reduce in complexity, aspects of the phenomenon that we find less interesting in order to examine more easily aspects we do wish to investigate' (Cook (2002), p.236). Thus, in building a model we must identify the aspects of the phenomenon we are interested in, and then simplify as much as it is required those aspects that are of less interest.

The second thing to notice is that, on this view, every model has two kinds of elements: *artifacts* and *representors*. The distinction between these two elements is just what you would expect: 'Call those aspects of the model that are intended to correspond to real aspects of the phenomenon being modeled representors, and those that are not intended to so correspond artifacts.' (Cook (2002), p.237). Presumably, the artifacts correspond to those aspects of the phenomenon we idealize or plainly misrepresent (sometimes correctly), and the representors to those aspects that we do not idealize. Now, the challenge this theory faces is to explain how to choose artifacts and representors in a way that isn't detrimental to one's theory.

In order to have a feel for what this kind of theorist has in mind, it is useful to start with a simple example:

> Little Victoria: Our objective is to construct a scale model of the famous Spanish ship, Victoria. To be more specific, what we want is a scale model of Victoria's external shape. We want our model to be as accurate as possible, within reason; we don't care about differences at the molecular level or differences that are due to the materials we use. Our ship, call it 'Little Victoria', isn't an exact scale duplicate of Victoria. There are many differences between these two ships. Little Victoria's internal makeup looks nothing like Victoria's. Where Victoria has chambers and bedrooms, Little Victoria has many sticks going from side to side. The sole purpose of these sticks is to help the model hang together. Despite the many internal differences between these two ships, they are exceptionally alike in their external appearance—setting size and materials used aside. Given our objective, it is plausible to say that Little Victoria is, in fact, a successful scale model.

It is easy to apply the artifact/representor distinction to this example. The internal sticks in Little Victoria are clearly artifacts of the scale model. Similarly, all the ways in which Little Victoria differs from Victoria can be taken to be artifacts of the scale model, so long as these differences do not translate into differences in Little Victorias external shape. What are the representors? Well, all the aspects of the scale model that are directly relevant to the external shape of the model are taken to be representors. Thus, the shapes of each individual external part, and the proportions of each part with respect to the rest are taken to be representors. Presumably, we can say that Little Vicotria is a successful scale model, because its represantors line up nicely with the aspects of Victoria that we want to model and the artifacts do not get in the way of doing so—they only misrepresent aspects we don't care about.

An key thing to notice, and that will play a crucial role later, is that a correct choice of artifacts and representors is not independent from the goals of our model; certain goals allow for artifacts that others won't. If our goal were to build an exact scale Victoria replica, Little Victoria wouldn't make the cut. If Little Victoria were part of an exact scale replica contest, we wouldn't persuade the judges by saying that the random sticks inside our model are just artifacts and that, consequently, we shouldn't be penalized. Relative to this goal, we cannot reasonably take those sticks to be artifacts of the model; they are undesirable features that need to be eliminated. If Little Victoria is to be an exact scale replica, it better be that her interiors look a lot like Victoria's. For this new goal we need fewer artifacts—although it isn't altogether clear to me which ones we can preserve and which ones we can't. Thus, the correct choice of artifacts and representors is goal dependent.

Now that we have a better grasp of the artifact/representor distinction, the question we should attend is: how to choose artifacts and representors for a theory of vagueness? Different theorists of vagueness would answer this question in a different way. For the sake of simplicity, let's consider what a degree theorist (Machina (1976), Edgington (1997), Smith (2008)) may say—our criticism can be used against other theories as well.

Presumably she would say that the existence of degrees of truth and the logical relations between sentences predicted by her model should be taken to be representors. Accordingly, she would like to say that the particular assignment of real numbers to sentences, and all the cut-offs and precision that come with such an assignment, ought to be regarded as artifacts. Thus, given this choice of artifacts and representors, what this theory cares about is the existence of degrees of truth and the logical relations obtaining between sentences; the rest are just artifacts. If this is the right way of thinking about this issue, then it would be unreasonable to complain that this theory misrepresents vagueness by making it precise. This seems to be exactly what Cook (2002) has in mind:

> We are misdescribing our linguistic practice only if we assert that the sharp cut-offs provided by assignments of real numbers to sentences represent real qualities of the phenomenon[...]The problematic parts of the account are not intended actually to describe anything occurring in the phenomenon in the first place, then they certainly cannot be misdescribing.(Cook (2002), p.237)

Part of the view is, then, that only representors can misdescribe and, given that all the cut-offs are merely artifacts, degree theory isn't misrepresenting vagueness. If degree theorist has chosen their artifacts and representors correctly, then this line of thought seems to be in order.

Now, the question is whether this is the right way of thinking about this issue. I shall argue that it isn't. As we have pointed out, the choice of artifacts and representors is restricted—in good part—by the goals of one's theory. Artifacts are intended as an aid to achieve one's goal, and it is indispensable that they play that role. What, then, is the objective of a theory of vagueness? A good way of finding the answer to this question is by paying attention to what has been obsessing philosophers for over four decades. What is it that we find so incredibly puzzling? The first thing that comes to mind is the Sorites Paradox. However, a moment of reflection shows that this paradox is only a symptom of something that goes deeper. What motivates the Sorites Paradox to begin with is the thought that vague predicates classify without setting sharp boundaries, without drawing any cut-offs. This is exactly what vagueness is all about, and what philosophers have strived to explain for a long time—notice that even the epistemicist wants to capture this at the epistemic level. As such, the main objective of a theory of vagueness ought to be to explain how vague predicates can classify without drawing sharp cut-offs. This theory may have other objectives as well—like an account of the logical relations between vague sentences, and a solution to the Sorites Paradox. As such, the choice of artifacts must be guided by this goal and perhaps others as well.

If this is correct, and it is, we cannot just wave our hands and call the uncomfortable cut-offs artifacts of the model. Doing so is tantamount to having a theory of motion that doesn't allow for movement. One cannot simply say: 'In my theory of movement I idealize space away. The theory is much simpler this way'. A theory of vagueness cannot afford to have cut-offs dividing two different semantic categories as artifacts, or as anything else. Vagueness is a very delicate phenomena; a simple cut-off in the wrong place is enough to vanish it. Now, one's theory can have as an objective the goal of explaining logical relations between sentences of a language admitting several degrees of truth. This is a legitimate theory, with a legitimate goal. However, this isn't a theory of vagueness—it is insufficient to quench the philosophical worries that the phenomenon gives raise to. The bottom line is this: if one's objective isn't to model vagueness, then it is a good idea to draw artificial cut-offs, otherwise doing so is out of the question.

We find ourselves in an extremely difficult position. There is one prime requirement that theories of vagueness must respect, but it is incredibly hard to see how one could possibly fulfill that requirement. One response to this conundrum is to take this requirement to be unrealistic, and allow a theory of vagueness that doesn't satisfy it. This is precisely the option that Edgington (1997) recommends:

> The demand for an exact account of a vague phenomenon is unrealistic. The demand for an account which is precise enough to exhibit its important and puzzling features is not. I do not deny that there is higher-order vagueness—that a sorites can be run on 'clearly red'. But I am urging that we can get a good enough understanding of what is going on in a sorites series on 'red', while ignoring higher-order vagueness. (p.308/309)

As we have seen, I do not recommend this option. I don't think Edginton's theory—brilliant as it is—helps to explain any puzzling features of the phenomenon. The bare existence of borderline cases of various degrees is hardly puzzling, and that is what this theory can explain—along with logical relations between sentences in a language that admits borderline cases. What is deeply puzzling is that predicates can classify without sharp boundaries, and what we are lacking is a theory that at the very least is compatible with this.

In a sense I agree with Edgington's claim that an exact account of a vague phenomenon is unrealistic. Here is what I think is impossible to do: one cannot give an exhaustive description of the range of application of a vague predicate relative to a Sorites series. In other words, it is impossible to specify for each member of a Sorites series its status with respect to the relevant vague predicate. If there is something we can learn from all these theories is that any attempt to do this leads to unacceptable cut-offs. On my view, what philosophers have failed to recognize is that the phenomenon of vagueness resists to be described completely; one of the prime characteristics of this phenomenon is that there is something that it makes impossible to describe. What is this thing? The answer seems to be rather clear: the complete range of application of vague predicates.

On the face of this one can take Edginton's route, which leads to incompatibility with the phenomenon of vagueness, or one can build a theory that respects the fact that this phenomenon makes it impossible to fully describe the range of application of vague predicates. This theory won't be ideal, since it has to be silent about many things. However, it is important to recognize that this is the best thing that we can do—vagueness won't allow for more. Trying to do more than this is to go beyond the limits of what can be explained. But what exactly does a theory of this kind would look like? The alternative I recommend is *Semantic Quietism*.[22]

[22]For their helpful comments, I'm grateful to Bob Beddor, Marco Dees, Andy Egan, Ernie Lepore, Vann McGee, Brian McLaughlin, Lisa Miracchi, Agustín Rayo, Brian Weatherson, and Tobias Wilsch.

Chapter 3

The Quietist's Gambit

3.1 The Phenomenon

Predicates are tools of undeniable importance. We use them to draw classifications (i.e. we use the predicate 'x is food' to classify some things as food, and others as something else). This makes it possible to communicate and to speak truthfully. Understanding classifications doesn't seem to be particularly challenging, until we take vagueness into account. Then any confidence we may have about this subject goes out the window. Vague classification have proven to be extremely puzzling—as a result, vague predicates become quite problematic. Let's take a moment to appreciate some of these puzzling features. We shall proceed by way of an example.

Example 1: Good Runners

> You are observing the leading runners of today's 5k. They are very fast and in excellent shape. Of course, 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 are slow and out of shape. This is a nice sorites series. A friend approaches you and asks: 'did you have a chance to see good runners?' To which you reply: 'Yes, the fast ones are good runners.'

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's also clear that based on your assertion you didn't classify some members as good runners; the last ones have not been classified as a good runners. Thus, you have used ‘fast’—and ‘good runner’—to classify some members of the series in a certain way and not others. This much is uncontroversial, or at least it should be.

Now, it is also clear that based on your assertion we cannot find any good reason that would help us identify *the* last member of the series that has been classified as a good runner. The default position is that this is so because there is nothing like the last member of the series that is *determinately* a good runner.[1] The existence of such a member would entail that there is a pair such that the first member is determinately a good runner, but not the second, even though their running abilities are indistinguishable for all practical purposes. The existence of such a pair seems, *prima facie*, absurd.

This kind of classification has two central features:[2]

a. Relative to a suitable sorites series vague classifications have some positive and some negative cases.

b. There is no last positive case and there is no first negative case—assuming that the series is arranged from positive to negative.

Let’s take a moment to assimilate how odd this kind of phenomenon is. There are classifications with positive cases at one end and negative cases at the other end. However, there is no point at which the positive cases stop and no point at which the negative cases begin. Given this, how can there be a transition from the positive to the negative case? If there is no point at which the positive cases end, how can the negative cases come to be? There is a paradox in sight.[3] The nature of vague classifications is extraordinarily odd, and yet, we use them in a very familiar way. A theory of vagueness is called for.

---

[1]Of course, Epistemicists are the notable exception.

[2]Admittedly, vagueness comes in many shapes and forms. See Weatherson (2010) for arguments in favor of this claim. In this paper we will only focus on paradigmatic examples of vague classifications. Most of what we say here, if not all, can be said of other instances of vague classifications.

[3]A discussion of the sorites paradox can be found in section 8 and Appendix 2.

Before we move on, it is important to make explicit the big picture dialectic this paper presupposes. It could very well be that, relative to every sorites series, our linguistic practices and the way the world is fully determine a sharp cut-off between the positive and the negative cases at a *determinate* location. Thus, for all we know, Epistemicism may be true.[4] Yet, it is important to recognize that, at present, we have no clear idea of how this could be so. Such a sharp division has to be such that someone who is, say, 1.79m is tall, but someone who is a nanometer (a billionth of a meter) shorter isn't tall. We may think that if something determines such a cut-off, it has to be our linguistic practices in conjunction with the way the world is. However, we are clueless about how these two elements could deliver such a sharp cut-off, *rather* than some other sharp cut-off within a nanometer of difference.

Given this, we should take very seriously the hypothesis that vague classifications are sharp cut-off free. Thus, in what follows we shall operate under the assumption that Epistemicism is false and that there are vague classifications of the kind that has been described. I don't intend to be dismissive towards this view; on the contrary, I have tremendous respect for it. However, pending an important discovery regarding the determination of sharp classifications by way of using vague predicates, it is good practice to have serious alternatives.

3.2 Quietism

A theory of vagueness is, primarily, a theory of vague classifications. Logic, semantics, and pragmatics of vague languages are *particularly* interesting insofar as they are strongly related to them—some platitudes are worth stating.[5] This is why our starting point is vague classifications. Once we have learned how to think about them, we can move on to theorize about vague languages, their semantics, pragmatics, and logic.

To a first approximation, the view I wish to advance is that vague classifications

---

[4] See Williamson (1994) and Sorensen (2001).

[5] Epistemicism (Williamson (1994) and Sorensen (2001)), Nihilism (Unger (1979)) and some forms of Contextualism (Fara Graff (2000)), Gómez-Torrente (2010), Gaifman (2002), and Rayo (2008)) are very interesting theories, even if, in an important sense, they deny the existence of vague classifications. Their interest, however, derives from a different place.

are, in some sense, impenetrable: they resist full description. A classification can be *fully described* when it is possible to specify, for each relevant object, whether it is in the classification, outside the classification, or whether it has some sort of intermediate status.[6] Thus, a core feature of vague classifications is that they *cannot* be fully described. The view isn't that, due to our cognitive limitations, we cannot offer a full description of vague classifications. Rather, the view is that, as a matter of fact, vague classifications *are* such that they cannot be fully described. If there is a God, this is one of her limitations.[7]

This view gets motivated by a simple observation: a *full description* of a vague classification has to identify the last member of the relevant series that is in that classification. To be more specific, such a description has to identify all the members of the series that are, let's say, blond, and the first member that has some sort of intermediate (or negative) status. This would entail that there is a pair of members such that their hair color is *indistinguishable* and that one of them is flat-out blond and the other one something less than flat-out blond. Given that this is impossible, no full description can be given.[8].

This characteristic of vague classifications restricts the ways in which we can represent and think about them. Any formal tool we might use to represent this kind of classification has to do justice to the fact that no full description is possible. Therefore, our formal representation of a vague classification has to be absolutely unspecific about whether some relevant objects are in the classification, outside of it, or whether they have some sort of intermediate status. As I shall argue, and as Crysippus argued before,

---

[6]For example, consider the natural numbers from 1 to 10. The classification containing all even numbers relative to this group can be fully described. The numbers 2, 4, 6, 8, and 10 are in the classification, and the rest aren't.

[7]This is point of similarity between my view and an Epistemicism of the kind defended by Sorensen (2001). On both views the classifications we draw by using vague predicates are, to some extent, epistemically impenetrable. However, we think this for very different reasons. Thanks to Brian Weatherson for pointing this out.

[8]This is a simplified version of a very complex argument. A full version of this argument needs to take into account the Higher-Order Vagueness literature. This literature is highly technical, and for that reason I rather leave it out of this paper (except for a few remarks in section 4). A reader interested on a detailed discussion on Higher-Order Vagueness, and why I think that this simple argument is, at heart, correct, should read Chapter 2. A reader interested in less technical arguments of this kind that are, at heart, correct, can read Sainsbury (1996).

the proper way of being unspecific is to be *silent*[9]. Hence, the view I wish to advance is a variety of 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 description of vague classifications. What the view gains is compatibility with the phenomenon of vagueness; this move guarantees that our representation of vague classifications does justice to the phenomenon.[10] This is the Quietist's gambit. One of the major themes in this paper is that any view that aims for complete specificity has to sacrifice compatibility with the phenomenon of vagueness. As such, the Quietist's gambit is the best opening available.

So far I have only offered a sketch of a theory. In the following pages I shall explain the way in which Quietism departs from tradition—it directly rejects one of tradition's most entrenched dogmas—and why it needs to do so. This discussion serves a dual purpose: it helps to position the view in logical space, and it motivates the view, by showing why it is an improvement over tradition. After doing so, the task will be to turn this sketch into a painting, by providing details about the view and its utility. In particular we shall spend some time developing a semantic theory that is friendly to Quietistic inclinations. Quietism needs such a theory; without it its plausibility becomes dubious.[11]

3.3 Frege's Paradise

There is a long an venerable tradition according to which semantic reality is absolutely crisp and can be described, at least in principle, with the utmost completeness and precision. This tradition can be traced back—at least—to Frege (1893). In *Grundgesetze* §56 he is quite explicit about this:

[9]Brian McLaughlin (p.c) recently pointed out to me that Crysippus, the Stoic, had a similar view. A significant difference between our views is that Crysippus's view is Epistemicist at heart.

[10]Section 5 explains with all detail why this is so.

[11]As it turns out, the semantic theory developed in this paper can be used by Quietists and non-Quietists alike. Any theorist that recognizes that none of the existing theories of vagueness does a good job at modeling vague classifications can find useful the tools that this paper develops. Thanks to Agustín Rayo for pointing this out to me.

> A definition of a concept (of a possible predicate) must be complete; it must unambiguously determine, as regards any object, whether or not it falls under the concept (whether or not the predicate is truly ascribable to it). Thus there must not be any object as regards which the definition leaves in doubt whether it falls under the concept; though for us human beings with our defective knowledge, the question may not always be decidable. We may express this metaphorically as follows: the concept must have a sharp boundary...a concept that is not sharply defined is wrongly termed a concept. (p.259)

Let's call this kind of picture *Frege's Paradise*—a world where semantic reality, and our description of it, is complete and crisp.

Of course, when Frege wrote this passage, he had in mind a formal language powerful enough to derive most of mathematics; that is why absolute precision is not negotiable.[12] However, Frege's Paradise—under one guise or another—has had tremendous impact in philosophy of language and semantics, even though Frege was quite explicit about natural languages not being admitted in it (Frege (1879) and Frege (1988)).[13]

There are two basic assumptions in paradise that I wish to focus on:

- **Bivalance**: For each object o in the domain and for each predicate ⌜F⌝ in the language, ⌜F⌝ is either true or false (where o is the referent of ⌜a⌝).
- **Exhaustiveness**: For each object o in the domain and for each predicate ⌜Fa⌝ in the language, the semantic status of o with respect to ⌜F⌝ is fully specified; either ⌜F⌝ applies to o, or it doesn't, or this is, to some degree, indeterminate.

Bivalence entails Exhaustiveness, but not the other way around. A semantics where there are more than two truth-values can satisfy Exhaustiveness, but not Bivelance—assuming that at least one sentence takes a value other than true or false. Bivalence is

[12]Frege took this idea seriously to the point that the Julius Cesar problem became a major concern in his project. See Frege (1988), and Heck (1997).

[13]To some extent, Frege was wrong about this, as the work of Richard Montague and his followers makes clear.

a particular way in which Exhaustiveness can be implemented. Exhaustiveness is the requirement that we should represent the relation between predicates and the world in a crisp and complete way; nothing is left to guesswork. For Frege, Bivalence is indispensable, so Exhaustiveness has to be satisfied.

The main reason why Frege's Paradise has been so influential is because it is extremely fruitful and elegant—the advances in logic and semantics that this framework made possible are astonishing. A language in Frege's Paradise can be studied with the kind of systematization and rigor that analytical minds are so fond of. Of course, theorists have found powerful reasons to modify and amplify Frege's original framework—a proper treatment of modality, indexicality, discourse dynamics, anaphora, indeterminacy, and so on, require machinery that isn't present in Frege's *Grundgesetze*.[14] Even though the framework has been modified, theorists, for the most part, haven't left Frege's Paradise; they have only brought in new furniture. In particular, as we shall see, Exhaustiveness is a basic assumption that hasn't been dropped.

3.4 Fall from Paradise

Not every linguistic phenomenon can be studied in paradise. Vagueness is an exception. There have been many attempts to theorize about vagueness in some corner or other of Frege's Paradise. Typically theorists of vagueness abandon Bivalence, but never Exhaustiveness. It is because they preserve Exhaustiveness that they are admitted in paradise. It is because they jettison Bivalence that they are confined to a corner. What we need is to forget about Exhaustiveness, at least when we take seriously certain features of natural languages.

In this section I shall argue that leading theories of vagueness assume Exhaustiveness. I shall also provide good reasons to think that Exhaustiveness is the root of serious problems in these theories. The outcome is that we need a new way of thinking about predicates—Frege's Paradise wasn't meant to be.[15]

---

[14] For example, see Kripke (1963), Kaplan (1989), Heim (1983), and Kamp (1975) .

[15] Of course, I don't mean to suggest that we should stop using this paradigm across the board; there is no better way of doing set theory, programming, formal logic, and so on. One important thing to keep

Let's focus on a standard sorites series for 'tall'. The first member is clearly tall and the last one is clearly short. It's easy to see how standard theories of vagueness assume Exhaustiveness. According to Three-Value Theory, *each* member of this series is either tall, not tall, or borderline tall. Degree Theory holds that *each* member of this series is tall to some degree or other. Supervaluationism has it that *each* member of the series is either tall in all admissible interpretations, or not tall in all of them, or tall in some and not in others. This list isn't complete, but it is enough to give us the flavor of what we are after; other theories are variations and complications over one ore more of these three themes.

Each of these theories assumes Exhaustiveness; they have earned a spot in Frege's Paradise. Also, each is committed to the following: there can be two subjects such that one is only a nanometer taller than the other and one of them is clearly tall, whereas the other is something a bit less than clearly tall. If we are willing to reject Epistemicism on the grounds that there is nothing in our linguistic practice and the way the world is that could determine a sharp division—down to the nanometer— between the tall and the not tall, we should equally reject these theories on the grounds that there is nothing in our linguistic practices and the way the world is that could determine a sharp division—down to the nanometer—between the determinately tall and the less than determinately tall. If we are to commit to these sharp divisions, we might as well commit to Epistemicism.

The problem with these theories *isn't* that they are committed to borderline cases, or degrees of truth, or to a notion of supertruth. As far as our arguments go, there is nothing wrong with the core features of these views. The problem is Exhaustiveness; it is because these theories assume it that they end up with sharp divisions where vagueness is supposed to be. In a sense, Exhaustiveness collapses theories of vagueness into varieties of Epistemicism, loosely construed. This is so because Exhaustiveness is a sharp cut-off generator, and, presumably, we cannot know where those cut-offs are.

---

in mind is that when one is abstracting vagueness away, it is a very good idea to assume Exhaustiveness. It is only when one takes seriously certain features of natural language that we are required to abandon this assumption.

Of course, things aren't this simple. An Exhaustiveness defender may argue that the kind of criticism we have been entertaining overlooks the phenomenon of Higher-Order Vagueness. 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. Thus, if $DTall(Maria)$ then we say that Maria is determinately tall. If $D\neg Tall(Maria)$ we say that Maria is determinately not tall. And if $\neg DTall(Maria)$ and $\neg D\neg Tall(Maria)$ we say that Maria is borderline tall. Exhaustiveness forces us to say that there is a pair of adjacent members of the series, let's say Olivia and Hugo, such that $DTall(Olivia)$ and $\neg DTall(Hugo) \wedge \neg D\neg Tall(Hugo)$. This is so even if the difference in height between Olivia and Hugo is incredibly small. This reeks like an unwanted sharp cut-off.

However, the argument continues, there is no sharp cut-off between the determinate cases and the borderline cases (between Olivia and Hugo) because even though Olivia satisfies $DTall(x)$, she also satisfies $\neg DDTall(x) \wedge \neg D\neg DTall(x)$. Similarly, even if Hugo satisfies $\neg DTall(x)$, he also satisfies $\neg DDTall(x) \wedge \neg D\neg DTall(x)$. Thus, the argument goes, the cut-off isn't determinately there; neither Olivia is determinately determinately tall, nor Hugo is determinately borderline tall. The strategy is, then, to move up one order, look down, and see that the cut-off isn't determinately there. The idea is to keep doing this until we are satisfied.

Higher-Order Vagueness is a very technical subject with an abundant literature (Zardini (2013), Fara (2003), Wright (1987), Gómez-Torrente (2002)). Rather than exploring the technical subtleties surrounding this topic, I will offer an informal presentation of the main issues. It is natural to think that if a theory *needs* to appeal to Higher-Order Vagueness in order to defend the plausibility of the first-order theory, it is probably because it got first-order vagueness wrong. Theorists typically appeal to Higher-Order Vagueness when they attempt to cover the tracks of unwanted precision. When one points at the cut-off between the positive cases and the borderline cases, they direct you to what they call second-order borderline cases. When you point at the cut-off between the second order determinate cases and the second order borderline cases, they direct you to third-order borderline cases. This game can be iterated for a long, long time. The crucial thing to notice is this: if this kind of theory didn't posit a cut-off between the first-order determinate cases and the first-order borderline cases to begin with, Higher-Order Vagueness wouldn't be playing such a central theoretical role.

Even if this particular kind of appeal to Higher-Order Vagueness wasn't wrong headed to begin with, there are pressing questions that put a substantial amount of pressure on this kind of view. One issue is that, relative to a finite sorites series, there can only be a finite number of borderline cases.[16] Therefore, if we go high enough in the orders of vagueness, we can point to a cut-off between some determinate cases and some borderline cases, without being able to hide this under any higher-order-borderline-region-rug. Thus, if we go up one more order, we can see that the cut-off is determinately there. Here is another platitude worth stating: a sharp cut-off is a sharp cut-off, whether it is at the first order or at a really high order. If we are suspicious of Epistemicism because it postulates a cut-off between the tall and the not tall, we should be equally suspicious of a theory that postulates a sharp cut-off at a really high order. If there is nothing in our linguistic practices and the way the world is that could determine the first cut-off, there is nothing in them that could determine the second one. Exhaustiveness breeds varieties of Epistemicism. Pending a massive and completely unexpected scientific/philosophical discovery that can reveal how these cut-offs get determined, we should consider seriously the hypothesis that these theories—in so far as they appeal to Exhaustiveness—are wrong headed.

3.5 The Sound of Silence

Vagueness in paradise is just like free will in a deterministic world; either you reject it, or you distort it to the point that it is no longer recognizable. Exhaustiveness is a very useful assumption that most theorists can afford. When one isn't trying to model vague classifications—and related phenomena—there is no harm in *supposing* that we can offer a full semantic profile of each object in the domain. However, as we have

[16]The stronger version of this result can be found in the second chapter of this dissertation.

seen, when our objective is to theorize about vague predicates and classifications, this assumption is untenable.

We need to learn how to work outside of paradise. The rest of this paper attempts to do that. According to Quietism, the main reason why Exhaustiveness doesn't hold is because it is impossible to offer a *full description* of vague classifications—it is impossible to say, for every relevant object, whether it is in the classification, outside of it, or whether it has an intermediate status. If vague classifications *are* like that, then any model of this phenomenon has to respect this fact. Given this, it seems that the only alternative open to us is to endorse a theory where the range of application of a vague predicate is only *incompletely specified*; there are some objects such that the theory has to be completely silent about their status with respect to the range of application of the predicate. Now, it is important to be clear about the notion of "incomplete specification" that we have in mind, and how this notion can be used to offer a genuine alternative to theories based on Exhaustiveness.

Let's start by using a few simple examples to elucidate the notion of *incomplete specification.* As it will become clear in a moment, this notion is independent from the existence of vague classifications, in the sense that one can offer partial specifications of the range of application of perfectly precise predicates. Thus, for the sake of simplicity, I will start by offering an incomplete specification of the range of application of a precise predicate, and later on I shall explain the sense in which this tool can be used to model vague classifications.

Let our domain be the natural numbers from 1 to 10. Now consider the predicate 'is an even number'. Here is a *full specification* of the range of application of this predicate relative to our domain:

- $[\![x \text{ is even }]\!] = \{2, 4, 6, 8, 10\}$
- $[\![x \text{ is not even }]\!] = \{1, 3, 5, 7, 9\}$

The sense in which this is a full specification is straightforward; for each object in the domain, it has been specified whether or not the predicate "is even" applies to it, or whether it is indeterminate.[17]

To a first approximation, this is how one among many *incomplete specifications* looks:

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

This isn't our official notation, but it's colorful enough to get across the main ideas—see Section 7 and Section 8 for rigorous alternatives. The comas with nothing in between represent the incomplete aspect of the specification. They leave open the possibility that there is something else in the classification, without committing to it—the statue analogy below helps elucidate this idea. Now, let's understand the sense in which this is an incomplete specification. This specification says something about how 'is even' classifies, but it doesn't say everything there is to say. For instance, this specification takes 2 and 4 into the positive extension of the predicate, and 1 and 3 into its negative extension, but nothing has being 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* specification. It is a fact that 6, 8, and 10 are in the range of application of "is even", however this incomplete specification is entirely silent about this. The specification doesn't say that 2 and 4 are the only members of the extension, and it doesn't say that 5, 6, 7, 8, 9, and 10 are negative or borderline cases. This specification is completely *silent* about the status of these numbers with respect of the range of application of 'is even'. It is in this sense that the specification is incomplete.[18]

Lets use an analogy in order to get a better handle on the notion of incomplete specification. Imagine you are a sculptor's model. She made a sculpture of you carrying a briefcase. This sculpture is very accurate in all respects, except that the sculpture is missing some parts; there is nothing corresponding to your lower torso and upper legs (Figure 1).

[17] Of course, in this particular case, there is no indeterminacy.

[18] Notice that an incomplete specification isn't some kind of metaphysically incomplete entity. Rather, it is an incomplete representation of a classical, indeterminate, or vague classification. Thanks to Andy Egan for helping me clarify this point.

(Figure 1. Sculpture by Bruno Catalano.)

This, I claim, is an *incomplete sculpture*. It represents correctly the upper and lower parts of your body, but it says nothing about the middle part of your body. Now, the sculpture doesn't represent you as not having middle body parts, nor does it represents you as having an indefinite middle body. Let's stipulate this wasn't the intention of the artist. Rather, the sculpture is completely silent about your middle body, and it is so in virtue of being incomplete.

Incomplete specifications, like incomplete sculptures, accurately capture some aspects of what they try to model. However, both incomplete specifications and sculptures remain *silent* about some aspects of what they represent; incomplete specifications are silent about the status of some members of the domain with respect to the relevant classifications and predicates, even if there is a fact of the matter about what their status is, and partial sculptures remain silent about some of the parts of the object they model, even if there is a fact of the matter about what those parts look like.

3.5.1 Incomplete Specifications and Vagueness

When one uses incomplete specifications to model vague classifications, one gets the following kind of picture: relative to a sorites series for 'tall', the first members get classified as tall, the last ones as not tall, and absolutely nothing is said about the middle range. The sense in which incomplete specifications violate Exhaustiveness is straightforward; in the silence region there are objects such that the theory is silent as to whether or not the relevant predicate applies, or whether this is indeterminate.

Now, if this is the view, a pressing worry suggests itself right away. When evaluating a theory of vagueness one of the first things one must check is whether the theory in question draws an unacceptable cut-off between the positive range of application of the relevant predicate and the rest. If it does, the theory becomes quite implausible; if it doesn't, then we are off to a good start. The presence of cut-offs is a source of concern because what is distinctive about vague classifications is that they don't draw cut-offs. Well, certainly, the view I'm suggesting draws two cut-offs; one between positive and silent cases, and another between negative and silent cases. Without further clarification this would be reason enough to dismiss the view. In what follows I shall argue that, after proper clarification, this shouldn't be a source of concern.

Not all cut-offs are harmful; it all depends on what they divide. Quietism draws harmless cut-offs, unlike views that assume Exhaustiveness. In order to appreciate this let's get an idea of what makes a cut-off a bad one. After doing this I shall argue that the cut-offs Quietism draws are not of that kind.

A bad cut-off is one that divides, relative to a sorites series, the positive (or negative) cases from all the rest. Thus, a cut-off of this kind divides the positive case and the borderline cases, or the positive cases and the negative cases, or the objects that are positive cases to degree 1, and the ones that are positive cases to degree .9, and so on. This kind of cut-off is bad, because what is distinctive about vague classification is precisely that they don't draw this kind of division. A theory of vagueness that postulates that kind of cut-off is, at best, suspicious. Here is an example of bad cut-offs:

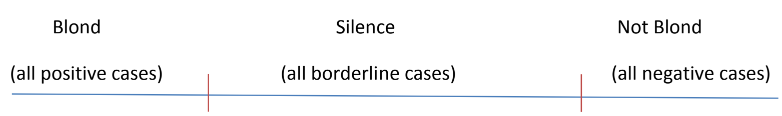

We need to draw cut-offs; that is the nature of our tools. We cannot draw them between all the positive cases and the first borderline case, because there is no such thing as *all* the positive cases (or *the* fist borderline case). Where else could we draw them? The first cut-off should be drawn between two positive cases and the second one between two negative cases. Thus, this is the picture I recommend:

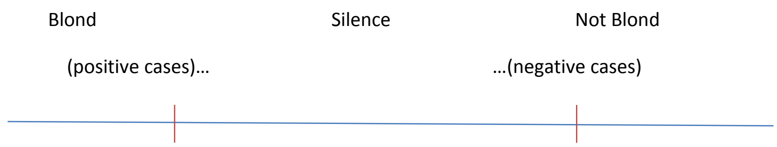

Notice that the first cut-off classifies most of the blond members of the series as blond, but not all of them. There is at least one member of the series who is blond, and our incomplete specification doesn't classify him as such—nor does the specification classify him as not blond, or borderline. The specification is simply silent about the status of that member of the series with respect to the range of application of the predicate. Similarly, this specification classifies most of the negative cases as not blond, but not all of them, since there is at least one member of the series who isn't blond, and the specification doesn't classify him as such.[19]

Why are these cut-offs not of the bad kind? The crucial thing to notice is that these cut-offs are perfectly compatible with the existence of vague classifications. This incomplete specification of the range of application of 'blond' doesn't misrepresent the relevant vague classification; it's just that it doesn't represent it completely. What makes this possible is, of course, that this incomplete specification becomes silent before it is forced to misrepresent the vague classification. This is precisely what the cut-off represents; it represents the point at which the specification becomes silent, which isn't the point at which the positive cases end and the borderline cases begin.[20] As such, this cut-off—and the one between negative and silent cases—is harmless. This is the

[19]There are some ways of interpreting Chryssipus, the Stoic, according to which he holds a very similar view. Thanks to Brian McLaughlin for pointing this out. It is nice to know that, deep down, I'm a little bit stoic.

[20]Thus, the region of silence cannot be identified with the borderline region. Such an identification goes against Quietism's core ideas.

correct way of using incomplete specifications.

Reflection shows that this view is quite natural. There is no such thing as all and only the blond members of the series, so a cut-off cannot lie there. If we try to specify, for each member of the series, what her status is with respect to the range of application of 'blond', it is unavoidable that we will end up drawing such a cut-off—this is what the literature on Higher-Order Vagueness teaches us. It seems, then, that the best way to model a vague classification is to use an incomplete specification in the way that has been suggested; we can only specify the range of application of a vague predicate with respect to some—but not all— of the members of the series. The only reasonable option is to let our incomplete specifications become silent before they misrepresent the nature of the phenomenon they attempt to model. What we find in the silence region—what humans and gods cannot describe—is an amorphous mixture of clear cases, indeterminacy, and mystery.

3.5.2 Borderline Cases

Before we move on we should say a word about the Quietist perspective on borderline cases. As we already pointed out, if a Quietist model is silent as to whether Marco is tall, that is not to say that the model takes Marco to be borderline tall. The silence region is not the borderline region. As it turns out, Quietism is quite flexible with respect to which conception of borderline cases we wish 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 & McLaughlin (1995), Tappenden (1993)). If we go this way we can model the rang 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 clear and some borderine cases. This way we guarantee that we make room for the existence of borderline cases and the phenomenon of vagueness. In section 8 I show how to do this within a Superquietist framework.

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 that they are F and it’s not knowable that they are not F. Moreover, a Quietist could think of borderline cases in terms 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 isn’t F. Finally, she could think that something is borderline ‘F’ just in case there is no fact of the matter as to whether it is F (Field (2003)). All this flexibility is possible because Quietism is, for and foremost, a theory designed to ensure compatibility with the phenomenon of vagueness, and not a theory designed to capture a specific conception of borderline cases.

Finally, it would be instructive to contrast Quietism and Agnosticism (Wright (2001) and Wright (2003)). At a certain level of abstraction these two views resemble each other.[21] 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*. We say that a proposition p presents a quandary for subject S just in case, a) S doesn’t 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 doesn’t know that it is possible to know whether or not p, and e) S doesn’t know that it is impossible to know whether or not p. 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 models. It is because of this that a Quietist can know that Maria is tall, and, yet, let her model be silent about it. 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 to be agnostic with respect to the law of excluded middle whereas some varieties of Quietism can accept such a law (see Section 8 for more details about this). I suspect that Quietism and

[21]Thanks to Brian Weatherson for pointing this out.

Agnosticism can play well with each other. However, it is worth keeping in mind that they have sharp differences at a basic level.

3.6 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 offer a *sub-optimal* model of vague predicates and classifications.[22] When a theorist models the range of application of a vague predicate, she has to remain *silent* about some positive and negative cases. For example, even if Olivia is clearly tall, a good Quietist model can be silent about it. Let's consider a model that is silent in this way. This model is *sub-optimal* for the simple reason that it is silent about certain facts; in a sense, a model that isn't silent about those facts is a better one, given that it describes the world in a more accurate way.[23]

This feature of Quietistic models may make us worry that there are some Higher-Order Vagueness problem in the vicinity. The worry can be stated as follows. Some sub-optimal models are better than others. In fact, some of them are *good enough*, whereas others aren't. Whether a Quietistic model is good enough depends on our theoretical goals. Assume that our main goal is to model conversations conducted using a vague language. 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 isn't a particularly good model. It isn't useful for that purpose. Now, the reasoning goes, a model that only classifies the two tallest members as tall, and is silent about the rest, is a better model than the first one, but still isn't good enough for our purposes. So the question is: which is *the* first Quietistic model that is good enough for our purposes? As soon as the Quietist identifies such a model, she will be stuck with horrible Higher-Order Vagueness problems. Why is that *the* first model that is good enough? Isn't this supposed to be vague? What is going on? Clearly, something has gone wrong.

[22] For now, we should understand by 'model' a theoretical structure that interprets a language, or a language fragment. For example, a model assigns referents to names and extensions to predicates.

[23] Thanks to Brian Weatherson for a discussion that lead to this section.

The problem is with the question: 'Which is *the* first Quietistic model that is good enough for our purposes?' This may seem like a natural question, but it is also more than that; it is an invitation to join Frege's Paradise. The Quietist must reject such an invitation. What the Quietist can do is identify *some* models that are good enough for our purposes, but she also has to be silent about others, even if some of them are also good enough (in Chapter 4 I provide details about how to make this kind decision when one is modeling a conversation). For example, she can say that the models that classify the first 100 members as tall and the last 100 as not tall are good enough for the purposes of modeling a specific conversation (or simply the range of application of 'tall'), and be silent as to whether a model that only classifies the first 99 members as tall is good enough (even if, as things stand, that model is good enough). This is another way in which Quietism avoids Higher-Order Vagueness worries.

One may try to keep pressing the Quietist and point out that she decided to become silent at a particular point, but that there are other points where she could have become silent. For example, she could have said that the model that classifies the first 99 members of the series as tall is also a good model, or she could have been silent as to whether the model that classifies the first 100 is a good one. Thus, one could ask the Quietist to specify *all* the other ways we could have classified the good models. This, again, is an invitation to join Frege's Paradise. The Quietist response is the expected one. She can specify *some* of the other ways she could have classified the good models while remaining silent about the rest. After all, it isn't possible to do better than this, so it is a mistake to ask for more.

We could, of course, point out that the Quietist has made another semi-arbitrary decision, and we can ask her to specify all the other ways in which she could have become silent.[24] An alert Quietist will stop playing this game fairly quickly. If she keeps responding we would force her to make a mistake. The best thing she can do is to become silent all together. This isn't a problem. As we have seen, silence doesn't misrepresent anything. As such, the Quietist approach is compatible with the phenomenon

[24]Thanks to Brian Weatherson for pressing this point.

of vagueness. What matters is that she has been able to identify a range of models that are compatible with the phenomenon of vagueness *and* that are good enough for her theoretical purposes. We cannot blame the Quietist for not offering a full description of a vague classification; one can only attempt such a thing in Frege's Paradise, and the Quietist has no interest in that.

Now, if the Quietist can avoid the horrors of Higher-Order Vagueness problems by becoming silent, couldn't other views do exactly the same? They can't do it without embracing Quietism from the start.[25] Non-Quietist views cannot afford to become silent at some high order because the cut-offs they draw between determinate cases and other categories are incompatible with the phenomenon of vagueness; this is the upshot in section 4. That is why they need to keep marching up the orders in an attempt the hide the traces of unwanted precision. As we have seen, the cut-offs in a Quietist model are innocuous, so there is nothing to hide. This is why the Quietist doesn't need to march up the vagueness orders.[26]

There is another objection the Quietist has to respond to. It can be phrased as follows: Quietism isn't 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 isn't what we have been dreaming about. Furthermore, Quietism claims that we have been dreaming about an impossible theory; according to this view the best we can do is to offer a sub-optimal theory of vagueness (you can call it a non-theory of vagueness, if you like). This is, once again, the Quietist's gambit. The one thing we must keep in mind is that the vagueness phenomenon doesn't allow for more. Before dismissing Quietism on the basis of this objection one needs to hold in her hands an exhaustive theory that is compatible with the existence of

[25]Section 8 shows how a Supervaluationist can embrace Quietism from the start. Superquietism is the result. Of course, Superquietism is a variety of Quietism.

[26]It is worth pointing out that even if the Quietist were to define a silence operator S in the metalanguage on top of a determinacy operator D, she shouldn't worry about Higher-Order Vagueness problems. This is so because in a Quietist framework there can be a member of a sorites series for F that is $D^{\omega}F(x)$ immediately followed by one that is $SF(x)$ without a problem. Recall that a Quietist cut-off isn't a bad cut-off.

vague classifications. Until then, Quietism stands strong as our best alternative. The interesting question isn't whether Quietism qualifies as a theory of vagueness. Rather, the interesting question is how to theorize about semantics given that we can only do that in a sub-optimal way. The previous section contains part of the answer to this question. The remainder of this paper gets us closer to a complete answer.

3.7 Semantics, Logic, and Paradox

Quietism is a framework that can host different semantics and logics. Which semantics and logics we build on top of it has a major influence on how we approach central topics in the philosophy of vagueness, language, and logic. For example, which solution to the sorites paradox we endorse, as well as our treatment of penumbral connections, standards of comparison, the interaction between context shifts and vague expressions, among others, depend on this. Quietism, without further supplementation, is silent about these issues. The only thing that Quietism promises is that any implementation of the framework will be compatible with the phenomenon of vagueness, and, as such, won't have higher-order vagueness problems. To the best of my knowledge this is the only theory that can promise such a thing with a straight face.

Perhaps, the easiest way to illustrate this is by interpreting Quietism using a standard Three-Valued Semantics. Now, a theory of this kind usually represents the semantic values of vague predicates as a pair of non-mutually exclusive sets $\langle A^+, A^- \rangle$, such that A^+ represents the extension of the predicate and A^- the anti-extension. The borderline region consists in the collection of objects that is in neither of these sets. Now, the Quietist can interpret such a region as a silence region, rather than as the borderline region. Given this, it is trivial to understand how this framework can be used to deliver a Quietist logic. The resulting logic is just a Three-Valued Logic with a different interpretation. The substantial difference between a standard Three-Valued Semantics and its Quietistic interpretation is philosophical rather than formal. Where the standard Three-Valued theorist sees borderline cases, the Quietist sees a region we cannot fully describe, containing an amorphous mixture of clear cases, indeterminacy, and mystery.[27]

This particular way of implementing Quietism is completely silent about the status of the inductive premise in the sorites paradox. According to this way of thinking about things, the sorites paradox is an intractable phenomenon we cannot say much about. However, that is not to say that we cannot have substantial theories of certain aspects of vague languages and communication. This version of Quietism provides a framework where, even if the paradox is, in some sense, intractable, there is much we can say about other features of natural languages.[28]

There are, of course, other Quietist alternatives open to us. In Section 8 I develop Superquietism. This view is the fusion of Quietism and Supervaluationism. Superquietism, unlike Supervaluationism, is compatible with the phenomenon on vagueness; there are no nasty cut-offs to be found. Also, this view preserves classical logic for the D-free language fragment.[29] The core Superquietist intuition is that it is problematic to understand (super)truth in terms of quantification over *all* admissible (classical) 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. Of course, one must be careful when selecting the silence region, otherwise disaster will ensue (this is explained

[27]If we wanted, we could project this philosophical difference into the formal apparatus. For instance, we could represents the positive extension of vague predicates as a pair $\langle P, S \rangle$, such that P has *some* of the positive instances, and S contains the silence range. We could represent the negative extension as a pair $\langle N, S \rangle$ such that N has *some* of the negative cases, and S is the same as in $\langle P, S \rangle$(the exact same region of silence). According to this way of doing things, the reason why the region of silence is built into both the positive and negative extensions is because such a silence region contains some positive and some negative cases; however, the theory cannot say which ones and how many there are. This would be a way to capture the philosophical differences between Quietism and Three-Valued Semantics within this kind of framework. However, notice that this difference isn't a substantial formal difference, given that S can be easily defined as what's not in P and N (or as what's not in A^+ and A^-).

[28]Similarly, this version of Quietism is completely silent about some instances of penumbral connections. For example, if the model is silent about weather or not Olivia and Hector are tall, then it is also silent about whether 'if Olivia is tall, then Hector is tall' is true. This version of the theory can attempt to explain why we feel tempted to say that all penumbral connections are true by appealing to pragmatic principles (see Tappenden (1993) and Soames (1999) for examples of this kind of strategy).

[29]In order to get this result we need to assume local validity.

in detail in Section 8). By going this way Superquietism can theorize 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; some positve cases, some negative cases, and a region that cannot be described.

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 it's negation is always true. As it is explained in next section, Superquietism also validates all penumbral connections formulated in the object language.

3.8 Superquietism

If we build Supervaluationism into the Quietist framework, something we should expect to see is a region of *silence* between the positive, borderline, and negative cases. Without this it is dubious that we have made room for vagueness—this is the Quietist insight. Our job is, then, to figure out a way of modifying Supervaluationsim so that these regions of silence can emerge. Superquietism is the result.[30]

Admissible Interpretations

Our starting point is the notion of an *admissible interpretation.* In a standard Supervaluationist framework (Fine (1975), McGee & McLaughlin (1995), Keefe (2000)) the rage 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 the between 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

[30]Thanks to Brian McLaughlin for a helpful discussion regarding this section.

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 inadmissible 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. For the sake of simplicity, let's abstract away from these complications and let's focus on a language who's only predicate is 'Tall' and whose domain of discourse only contains a sorites series for this predicate. Let's call this series 'T' and let's assume that it has 100 members, where 1 is the shortest and 100 is the tallest. An admissible interpretation of this language—classical as it is—draws a cut-off somewhere in T. If an interpretation draws a cut-off 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 40–60. [33].

[31] It is worth mentioning that, strictly speaking, S_A isn't needed. 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 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.

[33] Thus, the interpretation that classifies all the members between 1 and 42 as not tall, and all the rest as tall, is an interpretation in P_A. Also, an interpretation that classifies all the members between 1 and 56 as not tall and the rest as tall is a member of P_A. However, an interpretation that classifies the members between 1 and 35 as not tall and the rest as tall, is not an interpretation in P_A—it draws

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 clearly not tall—19 isn't the first clearly not tall member of the series, since there is no such thing. 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.[34] 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—now idealized away—conversation. Our choice, however, wasn't completely arbitrary; no inadmissible interpretation could have been admitted in P_A and no admissible one could have been admitted in N_A.

Supertruth

We have taken our first step—we have made the range of admissible interpretations acceptable by Quietist standards. The second step is to find a way of defining *supertruth* and the *deterinacy* operator in the right way. Recall that the objective is to have a region of silence between the positive, borderline, and negative cases. Given that this

the cut-off in the wrong location. There are exactly twenty interpretations in P_A.

[34]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.)

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, there is no sense to be made of quantification over *all* the admissible interpretations. 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 surface. Once we get this right, the definition of the determinacy operator becomes obvious.

Semantic Clauses

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 & 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*.[35] The base model is a non-classical model of the vague language with a partial interpretation function—it is easy to interpret a partial function in a Quietist way.

[35]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.

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 chose 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 classical interpretation, and σ is a variable assignment):[36]

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.[37]

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, both P_A and S_A are built into the sufficient condition, but in the clauses 2-4 only either P_A or N_A appears in the sufficient condition. Notice, however, that these sufficient conditions are unquestionably right 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. Let's go back to our example once again, but this time let's spell it out in more detail. Our domain only contains the members of a sorites series for 'tall'. The first member of the series, 1, is the shortest, and the last member, 100, is the tallest. 'tall' is the only predicate in this language. Besides this predicate, the only

---

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

[37]Notice that, given that our metalanguage is classical, $\neg D\Phi$ is true in M, σ, iff Φ is not supertrue in M, σ.

other non-logical expressions are names for each member of the series; the name for the first one is '1', for the second '2', and so on. Let's call the interpretation that draws a cut-off between 43 and 44 'M_{43-44}'. Thus, M_{43-44} classifies all the members between 1 and 43 as not tall, and all the members between 44 and 100 as tall. In general, M_{n-m} is the interpretation that draws a cut-off between n and m, where m is the successor of n. Also notice that for every M_{n-m}, n is the last not tall member in that interpretation and m is the first tall member in that interpretation. With this notation in hand we can specify the members of P_A, S_A, and N_A as follows (where m is the successor of n):

- $P_A = \{M_{n-m} | n = 40$, or $n = 41$, or, ..., or $n = 59\}$
- $S_A = \{M_{n-m} | n = 20$, or $n = 21$, or, ..., or $n = 39$, or $n = 60$, or $n = 61$, or , ..., $n = 79\}$
- $N_A = \{M_{n-m} | M_{n-m} \notin P_A \cup S_A\}$

Let's proceed to identify the positive, borderline, negative, and silences regions.

Positive Cases ($DTall(x)$)

Given (1), it is supertrue, for each member of the series between 80 and 100, that they are tall. This is so because every interpretation in $P_A \cup S_A$ draws the cut-off somewhere between 20 and 80. Thus, there is no interpretation in $P_A \cup S_A$ that classifies 80 as not tall, so it is true in all the interpretations in $P_A \cup S_A$ that 80 is tall. Therefore, it is supertrue that 80 is tall. This follows from the fact that $M_{80-81} \notin P_A \cup S_A$—also notice that there is *no* $M_{n-m} \in P_A \cup S_A$ such that $n \geq 80$. The same can be said for every member of the series between 1 and 20. Therefore, given (5), $DTall(x)$ is true of every member of the series between 80 and 100. These are the positive cases.

Negative Cases ($D\neg Tall(x)$)

Following analogous reasoning we can conclude that the negative cases—the range of objects that satisfy $D\neg Tall(x)$—are the objects between 1 and 20. For instance, all the interpretations in $P_A \cup S_A$ are such that 20 is not tall. This follows from the fact that

$M_{19-20} \notin P_A \cup S_A$, and, in general, there is *no* $M_{n-m} \in P_A \cup S_A$ such that $n \leq 19$. Therefore, given (1), it is supertrue that it is not the case that 20 is tall (it is supertrue that $\neg$(20 is tall)). The same can be said of every member of the series between 1 and 20. Thus, given (5), $D\neg Tall(x)$ is true of all the members between 1 and 20. These are the negative cases.

Borderline Cases ($\neg DTall(x) \wedge \neg D\neg Tall(x)$)

Now let's see who the borderline cases are. By definition they are the objects that satisfy $\neg DTall(x) \wedge \neg D\neg Tall(x)$. Thus, by (2) and (5) the members of the series that satisfy $\neg DTall(x)$ are the ones that are not tall in at least one interpretation in P_A—notice that we are not quantifying over $P_A \cup S_A$ anymore. These are precisely the members between 1 and 59; this is so because there is a member of P_A that draws a cut-off between 59 and 60 (M_{40-41}) and there is no member of this set that draws a cut-off anywhere between 60 and 100. By (2) and (5), the members of the series that satisfy $\neg D\neg Tall(x)$ are the ones that are tall in at least one interpretation in P_A. These are precisely the members between 41 and 100; this is so because there is an interpretation in P_A that draws a cut-off between 40 and 41 (M_{40-41}), and there is no interpretation in this set that draws a cut-off anywhere between 40 and 1. Therefore, the borderline cases—the range of objects satisfying $\neg DTall(x) \wedge \neg D\neg Tall(x)$—are those between 41 and 59.

Silence

What we have achieved is this: the positive cases are between 80 and 100, the borderline cases are between 41 and 59, and the negative cases between are 1 and 20. The rest is *silence*. Our region of silence lies, then, between 21 and 40 (inclusive), and between 60 and 79 (inclusive). This is exactly what we wanted. For the purpose of illustration it may be useful to list these classifications:

- Positive Cases ($DTall(x)$): $\{80, 81, ..., 99, 100\}$
- Borderline Cases ($\neg DTall(x) \wedge \neg D\neg Tall(x)$): $\{41, 42, ..., 58, 59\}$

- Negative Cases ($D\neg Tall(x)$): $\{1, 2, ..., 19, 20\}$
- Silence: $\{21, 22, ..., 39, 40\} \cup \{60, 61, ..., 78, 79\}$

Now, it is important to notice that in this model we are not completely silent in the silence region; we do say, for example, that the objects in $\{21, 22, ..., 39, 40\}$ satisfy $\neg DTall(x)$, and that the objects in $\{60, 61, ..., 78, 79\}$ satisfy $\neg D\neg Tall(x)$. These are trivial claims that follow from how a sorites series is built. Importantly, we *are* silent about the substantial claims; we are silent about whether the objects in $\{21, 22, ..., 39, 40\}$ satisfy $D\neg Tall(x)$, and whether the objects in $\{60, 61, ..., 78, 79\}$ satisfy $DTall(x)$.

Superquietist Logic

It's time to say something about the notion of *validity*. Traditionally there are at least two options open to Supervaluationism; global and local validity.[38]

- Global Validity: $\Gamma \vDash_g \Phi$ iff if every $\Psi \in \Gamma$ is true in all admissible interpretations, then Φ is true in all admissible interpretations.
- Local Validity: $\Gamma \vDash_l \Phi$ 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.

Whether Supervaluationism should accept global or local validity has be subject to dispute. 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 hold to global validity, there are no problems that emerge from doing so. McGee & McLaughlin (1995) and Varzi (2007) argue that the right notion of validity is the local one.

I will let Supervaluationists resolve their own disputes. In what follows we shall see how Superquietism can adopt a version of local validity. We need to modify the

[38] For more alternatives see Varzi (2007).

definition of local validity if it is to be acceptable by Quietist standards. In particular, we have to avoid quantification over all admissible interpretations, because, as Quietists, we find such a notion quite troubling. Rather than quantifying over all admissible interpretations, we should quantify over all the interpretations in $P_A \cup S_A$. The resulting clause is:

- $\Gamma \vDash_l \Phi$ 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.[39]

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 $P_A \cup S_A$.

Another nice feature is that all penumbral connections are supertrue. 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 the interpretations in $P_A \cup S_A$, all penumbral connections are supertrue.

Let's consider a specific example. Take 47 and 48 in our example above. 48 is taller than 47. Thus, the following sentence is an instance of a penumbral connection:

(a) If 47 is tall, then 48 is tall.

As such, (a) is supertrue according to Superquietism. This can be shown by the fact that the only interpretations where 47 is tall are interpretations that draw the cut-off somewhere between 1 and 47.[40] But 48 is also tall in all those interpretations. Therefore, (a) is supertrue.

[39] Notice that, as expected, $\Phi \vDash_l D\Phi$ isn't locally valid.

[40] Some of these interpretations are not in $P_A \cup S_A$, but that doesn't matter in this case.

Similarly, the following instance of a penumbral connection is supertrue:

(b) It is not the case that 47 is tall and 48 is not tall.

The reason for this is exactly the same as the reason why (a) is supertrue. The only interpretations where 47 is tall are the ones that draw the cut-off somewhere between 1 and 47. But in all those interpretations 48 is tall as well, so it superfalse that 47 is tall and 48 is not. Therefore, it is supertrue that it is not the case that 47 is tall and 48 is not tall.

As it turns out, Superquietism can solve the Sorites Paradox in the same way Supervaluationism does. According to Superquietism, relative to a sorites series for F, (a) and (b) are both supertrue (where x' is the successor of x):

(a) $\neg\forall x(F(x) \to F(x'))$

(b) $\exists x(Fx \wedge \neg F(x'))$

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, (a) and (b) are true in all the interpretations in these set.

Superquietism isn't as pretty as Supervaluationism; we don't validate as many sentences as they do. That is the price we pay for our silence. For example, in a Supervaluationist framework the following sentence is always true:

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

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 61 in our example above. Our model is *silent* as to whether $DTall(61)$—61 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 $\neg DTall(61)$—there is no interpretation in P_A where 61 is not tall. Therefore, the model is silent as to whether $DTall(61) \vee \neg DTall(61)$. Similarly, our model is also silent about $DTall(61) \to DTall(61)$ for the exact same reasons.

It's worth noting that the model is also silent about 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 all these are acceptable prices to pay. Superquietism is compatible with the phenomenon of vagueness and preserves classical logic for the fragment of the language that is D-free. This is a significant improvement over other theories. Now, the fragment of the language that contains D isn't as well behaved as we may wish. However, keep in mind that D is only a metalinguistic operator; if we have to be silent about some penumbral connections involving it, we shouldn't worry that much, so long as we get the penumbral connections right in the object language. That we can do.[41]

[41] For their many helpful comments, I'm grateful to David Black, Peter van Elswyk, Andy Egan, Will Fleisher, Simon Goldstein, Dirk Kindermann, Ernie Lepore, Brian McLaughlin, Lisa Miracchi, Agustín Rayo, Carlotta Pavese, and especially Brian Weatherson.

Chapter 4

Silence in Conversation

4.1 Introduction

Philosophers and linguists have studied vague languages and linguistic communication quite extensively. However, relatively little attention has been given to the phenomenon of vague linguistic communication—that is to say, communication conducted using vague languages. Some progress has been made in this direction, but theories of vague linguistic communication remain fundamentally flawed. Their problem is that of compatibility; the models that these theories offer are simply incompatible with the existence of vague communication. Some theories (Stalnaker (1978), Manor (2006), Rayo (2008)) attempt to get around this problem by invoking pragmatic mechanisms, whereas others (Kamp (1981a), Barker (2002), Gaifman (2002), and Shapiro (2003)) try to do so by relying on heavy semantic machinery. On my view, a very different approach is required. This is not to say that these theories haven't made any progress, it is just to say that their foundations stand in need of improvement.

The main aim of this paper is to offer a theory of vague linguistic communication that is compatible with the phenomenon of vagueness and robust enough to model the evolution of conversations. In order to do so we need adequate theories of linguistic communication and vague languages. I shall adopt a theory of linguistic communication that is broadly Stalnakerian. However, an important departure from standard versions of this theory is required; we need to reject one of its basic principles. This has to be done in order to make room for the phenomenon of vagueness. The theory of vagueness I shall adopt is *Quietism*.[1] The main strength of this theory is that, unlike others,

[1]See Chapter 3 for an explanation of this view.

it guarantees compatibility with the phenomenon of vagueness—it is the only theory that doesn't idealize vagueness away. The strategy is to make a Stalnakerian theory Quietistically friendly. The result is a theory of vague linguistic communication that makes room for the phenomenon of vagueness and that is robust enough to model the evolution of conversations.

First I shall spend some time arguing that standard Stalnakerian theories (Stalnaker (1978), Barker (2002), and Rayo (2008) have serious difficulties modeling certain kinds of vague conversations. Then I shall argue that common ways of fixing this problem also face serious difficulties. Finally I show how a broadly Stalnakerian theory can be incorporated into a Quietist framework in order to guarantee compatibility with the phenomenon of vagueness. This, however, can only be done at a modicum price. The core idea is that if we want our Stalnakerian theory to be at the very least compatible with vague communication, then we must realize that it is in principle impossible to model every single aspect of the context set and the content of vague sentences.

To a first approximation, the view is that our Stalnakerian theory needs to be silent as to whether a) certain live options in a given conversation are in fact live options, and b) whether certain vague sentences are true relative to some possibilities, even if those sentences are in fact true (or false) relative to those possibilities. Hence, phrased in Stalnakerian jargon, the view amounts to this: there are certain robust aspects of the context set and content of vague sentences that our theories must be *silent* about. In particular, there are certain worlds that *are* in the context set but that our theory has to be silent as to whether this is so. Similarly, there are worlds that *are* part of the content of vague sentences, but our theory has to be silent as to whether this is so. If this is done right, then we can make a Stalnakerian theory of linguistic communication compatible with the possibility of vague linguistic communication. This is the Quietist promise.

4.2 The Stalnakerian Framework

According to the Stalnakerian picture (Stalnaker (1978)), linguistic communication takes place against a background of information. This body of information plays a dual role; it is the information relative to which speech acts are interpreted, and it is the information that speech acts have an effect on. We call this body of information the *common ground*, which is the set of presuppositions in play in a conversation at a particular time. A presupposition is a proposition that the participants to the conversation are willing to accept as if it was common knowledge. Thus, in conversations we often presuppose that our interlocutors speak English, that Obama is the president, that we live on planet Earth, that it's the 21st century, among many other things. As conversations evolve, and assertions take place, more propositions are added to the common ground, provided that the participants to the conversation accept what has been asserted.

The common ground determines the *context set*; the set of all the possibilities that are compatible with what is in the common ground. These worlds are the live options at that stage in the conversations; they represent the ways the world might be for all that has been presupposed. Thus, if in our conversation we presuppose that we all speak English, the context set only contains possibilities where we all speak English—all the possibilities where one of us doesn't speak English have been excluded. Similarly, if we presuppose that Obama is the president, all the possibilities in the context set are such that Obama is the president. Therefore, if we presuppose these two propositions, all the possibilities in the context set are such that we all speak English and Obama is the president. Now, if we don't presuppose anything in particular about the youngest U.S. chess master's age, then the context set contains possibilities where Clarissa Yip is quite young, others where she is a bit older, others where she is a bit older, and so on.[2] In short, the context set contains the possibilities that are still live options at a given stage in the conversation—it is a live possibility that Clarissa Yip is very young,

[2]Of course, in a normal conversation we would presuppose that she isn't a toddler, and that she isn't a hundred years old either. Vagueness is lurking.

but it isn't a live possibility that Obama isn't the president.

One of the objectives in a conversation is to eliminate some of the live possibilities from the context set and keep others; this is, in good part, what it is to exchange information in a conversation. This is usually achieved by engaging in the practice of *assertion*. Let's say that you want to know whether Clarissa Yip is very young—recall we are not presupposing anything in particular about her age at this point in the conversation. If I assert that Clarissa Yip is 9 years old, I thereby propose to eliminate the possibilities where she is older, and also the possibilities where she is younger.[3] If you accept my assertion, then we come to presuppose that Clarissa Yip is 9 years old; the effect of this is that the possibilities where she is older or younger are no longer live options.

Within this picture it is natural to take the contents of declarative sentences to be sets of possible worlds; the set of possible worlds where the sentence is true. Thus, the proposition expressed by 'Clarissa Yip is young' is the set of all the possible worlds where she is young. Given this, the following picture of how an assertion affects the context set looks very natural: an assertion of 'Clarissa Yip is young' is a proposal to intersect the proposition expressed by this sentence with the context set. If the proposal is accepted, what we get is a context set that only contains worlds where Clarissa Yip is young. This is a beautiful picture, and a very useful one; these two things don't hang out together very often.

As it stands, it isn't obvious how to make the Stalnakerian machinery vagueness friendly. Both context sets and propositions have been represented as classical sets of possible worlds. One thing we are clear about in the vagueness literature is that classical sets—just like that, without fancy footwork—repel vagueness. This is, of course, no reason to reject this framework; it is only a reason to refine it.

---

[3] I'm assuming that update is, to a large extent, a pragmatic matter (Stalnaker (1978)) rather than a semantic one (Groenendijk & Stokhof (1991)). However, nothing I say hangs on this. A Quietistic interpretation of dynamic semantics is outside the scope of this paper.

4.3 Pragmatic Evasion

Now, one may think that proper understanding of this machinery urges us not to worry too much about incorporating vagueness into it. After all, the story goes, there is a principle of rational communication that restricts the effects of vagueness in discourse. Recall the three principles of rational communication that Stalnaker (1978) advanced:

I A proposition asserted is always true in some but not all of the possible worlds in the context set.

II Any assertive utterance should express a proposition, relative to each possible world in the context set, and that proposition should have a truth-value in each possible world in the context set.

III The same proposition is expressed relative to each possible world in the context set.

Principles (I) and (III) are largely orthogonal to our current discussion. The principle that is directly relevant for us is (II). It is plausible to think that a vague assertion expresses a vague proposition—or a vague collection of precise propositions—that may be indeterminate in some of the possible worlds in the context set.[4] For example, if we don't assume anything in particular about Clarissa Yip's age, then there is a world in the context set where she is borderline young. If I were to assert 'Clarissa Yip is young' the proposition expressed by this sentence wouldn't be true or false relative to the possible worlds where she is borderline young. Principle (II) condemns assertions of this kind.

This principle holds some weight in the Stalnakerian framework, given that it serves as a bridge connecting semantic and pragmatic presupposition. Here is how the explanation goes. If p semantically presupposes q, and q isn't pragmatically presupposed—isn't

[4]Quietism is neutral as to whether a vague sentence express a vague proposition or a vague range of precise propositions . It is easier to run explanations under the assumption that sentences express vague propositions. For that reason I shall proceed as if this assumption was correct. However, everything I say can be easily translated into a view according to which sentences express a vague range of precise propositions (Lewis (1975)). Quietism can also incorporate supervaluationist intuitions. The result is a view that I call *Superquietism* (see Chapter 3).

in the common ground—an assertion of p will be neither true nor false in those worlds in the context set where $\neg q$ is true. Hence, given principle (II), an assertion of p, relative to this kind of context, triggers the phenomenon of proposition accommodation; the audience will accommodate by coming to presuppose—in the pragmatic sense—p's semantic presupposition (q in this case). Given the services that it provides, this principle has been cherished by some; Von Fintel (2008) goes as far as to call it '[an] irreducible property of natural language pragmatics.'

This principle is fairly clear—what we need to understand is Stalnaker's motivation for it. Here is the idea:

> The rationale for this rule is as follows: The point of an assertion is to reduce the context set in a certain determinate way. But if the proposition is not true or false at some possible world, then it would be unclear whether that possible world is to be included in the reduced set or not. So the intentions of the speaker will be unclear. (p.90)

The idea, I take it, is that vague assertions get in the way of successful communication, given that they don't allow for the context set to evolve in a determinate way. Thus, vague assertions may induce a conversational crash, bringing the conversation to a halt. This is certainly true in some cases. Let's consider an example where indeterminacy does get in the way of communication.

> **Valeria's House**: You want to get to Valeria's House and she is giving you directions over the phone. At some point she gets you to the right street, and both of you presuppose so. There are only two houses on that street: a bluish-gray and a bluish-purple house. Let's assume that both houses are borderline blue, and it isn't the case that one of them is bluer that the other. Valeria says: 'Knock on the door. My house is the blue one'.[5]

Principle (II) correctly condemns Valeria's assertion. Here is why: at this point in the conversation the context set only contains two kinds of worlds—worlds where her house

[5]This example is based on examples in Dorr (2003) and Rayo (2008).

is bluish-gray and worlds where it is bluish-purple. Now, Valeria's assertion of 'My house is the blue one' expresses a proposition that is, presumably, indeterminate in all the worlds in the context set—recall that both houses are borderline blue. Thus, there is no reasonable way to update the context set based on her assertion. As such, her assertion is clearly defective, and she need to take it back, clarify it, or expect some sort of pragmatic repair.

But how can there be successful communication if a large portion of our assertions contain vague vocabulary? Wouldn't principle (II) condemn must of what we say? Not really, if the context set cooperates! There are contexts where an assertion of 'Clarissa Yip is young', vague as it is, can be in perfect harmony with (II). Consider, for instance, a conversation where the only worlds in the context set are such that Clarissa Yip is either between 6-9 years old or between 70-75 years old. Relative to this context set an assertion of 'Clarissa Yip is young' is in line with (II). The reason being that this assertion is either true or false in every world in the context set—it is true in all the worlds where she is between 6 and 9 years old and false in all the worlds where she is between 70 and 75 years old. After such an assertion the communicative intentions of the speaker are perfectly clear and context set can evolve in a determinate way.

Valeria's House is an example where indeterminacy slaps us in the face; it is a case where the fact that the assertion is indeterminate relative to some (all) of the worlds in the context set does get in the way of successful communication. There are many cases such that, given the characteristics of the context set, indeterminacy isn't an issue. What I wish to argue now is that in a good number of well-conducted conversations, assertions can be indeterminate relative to some worlds in the context set without this having a major effect on the evolution of the conversation. As it turns out, these are cases where indeterminacy goes unnoticed for the most part; this is exactly what happens when vague expressions are used correctly. If this is right, then the plausibility of principle (II) should be put into question. Let's get into this by way of an example.

Fodor's Apartment: We are visiting New York City for the first time, and we have a long to-do list. The next item in our list is to visit Jerry Fodor's

> apartment and we have no clue about its location. Our tourist guide book isn't helpful so we decide to ask for directions. A friendly bystander asserts: 'Fodor's apartment is near the Metropolitan Opera House.' We know where the Metropolitan is, so we start moving in that direction.

This is, by any reasonable standards, a fine assertion. However, as we shall see, it is very hard to understand why this is so if we hold to principle (II).

Notice that prior to the helpful bystander's assertion our context set contained worlds where Fodor's apartment is by Central Park, others where it is in Harlem, others where it is Downtown, and many others. The reason is that we—helpful bystander included—didn't presuppose anything in particular about his apartment's location. In a sense, our context set is soritical with respect to the predicate 'x is near the Metropolitan Opera House'. This is so because in this context set there is a world where the apartment is right next to the Metropolitan, another world where it is just a bit farther away, another one where it is even a bit farther away, and so on, until we get to a world where the apartment is as far away as a NYC apartment can be from the Metropolitan. Now, consider the perfectly fine assertion of:

> (FA) Fodor's apartment is near the Metropolitan Opera House.

Certainly this sentence is true in worlds where the house is very close to the Metropolitan, and false in worlds where it is far away. However, it is reasonable to think that (FA) is neither true nor false relative to some worlds in the context set; think of worlds where the house is *more or less* near the Metropolitan. As such, principle (II) condemns this assertion. What is surprising is that this assertion seems to be perfectly felicitous; the conversation can continue smoothly after it. It seems, then, that there is something wrong with principle (II).

4.3.1 The Accommodation Strategy

One my try to rescue principle (II) by invoking some sort of pragmatic repair strategy—after all, as Rayo (2008) has shown, a standard Stalnakerian framework can help us understand many puzzling aspects of vague linguistic communication. The thought is that because—relative to our example—(FA) is condemned by (II), participants to the conversation, cooperative as they are, *accommodate* by modifying the context set in order to make the assertion compatible with this principle after all.[6] This strategy has been proposed by Rayo (2008)—although he proposed it in the context of defending a more general principle.[7]

The central idea is that accommodation is required in this case because an assertion that violates principle (II) is defective, and, as such, it is neither true nor false. Thus, assertions containing vague vocabulary impose some constrains on the common ground and context set. In order to avert defectiveness participants to the conversation accommodate by modify the context set in such a way that an assertion of (FA) is in line with principle (II)—and, thus, no longer defective. Hence, for example, participants in the conversation take the context set in our example and accommodate by transforming it into a context set where Fodor's apartment is either clearly close or clearly far away from the Metropolitan, thereby taking a soritical context set and transforming it into a *gappy* one. Presumably they do this by presupposing something that would eliminate worlds where the apartment isn't clearly colose or clearly far away from the Met. Relative to this updated context set (FA) is perfectly compatible with principle (II).

Let's define this terminology in order to have a better grasp on the view we are discussing. A series of worlds $\langle w_1, ..., w_n \rangle$ is *soritical* with respect to Fa just in case a is F in w_1 and not F in w_n and the degree to which a is F smoothly decreases by a tolerant margin from w_1 to w_n. A tolerant margin is one that doesn't affect the justice with which the predicate is applied.[8] A context set C is soritical with respect to a sentence $\ulcorner Fa \urcorner$ just in case C contains enough worlds to build a series of worlds that is soritical with respect to $\ulcorner Fa \urcorner$. In Fodor's Apartment the context set is clearly soritical with respect to (FA). A series of worlds $\langle w_1, ..., w_n \rangle$ is *gappy* with respect to sentence $\ulcorner Fa \urcorner$ just in case it is just like a soritical series with respect to $\ulcorner Fa \urcorner$ except that there

[6]Accommodation is understood in Lewis (1979) sense.

[7]He calls it *the principle of clarity*: "Make your assertions in such a way that it is clear to your audience how you propose that they be used to modify the conversation's presuppositions."

[8]I borrow this way of understanding tolerance from Wright (1976).

is an adjacent pair of worlds in the series $\langle w_m, ..., w_{m+1}\rangle$ such that $\ulcorner Fa \urcorner$ is true in w_m, false in w_{m+1} and the degree to which a is F in w_m is much larger than in w_{m+1}. A context set C is gappy with respect to a sentence $\ulcorner Fa \urcorner$ just in case the worlds in C can form a gappy series relative to $\ulcorner Fa \urcorner$, but not a sorites series of worlds with respect to $\ulcorner Fa \urcorner$. For example, a context set that only contains worlds where Fodor's apartment is either one block or sixty blocks away from the Metropolitan is one that is gappy with respect to (FA).[9]

Thus, an abstract description of the accommodation strategy is this. If a speaker asserts $\ulcorner Fa \urcorner$ in a context that is soritical relative to that sentence, participants in the conversation will tend to accommodate by adopting a presupposition that will change the context set into a gappy one. Once the context set is gappy with respect to $\ulcorner Fa \urcorner$ they can update the common ground with the content of the original assertion in a way that is compatible with principle (II). This is how, according to this view, we can explain Fodor's Apartment while preserving (II). The violation of the principle is only apparent, since the participants to the conversation accommodate the context set in order to make the assertion compatible with it after all.

4.3.2 Accommodation Won't Help

Accommodation happens, and it happens quite often. There is plenty of linguistic evidence to support this claim (Lewis (1979), Beaver & Zeevat (2004), Von Fintel (2008))[10]. Furthermore, there are good arguments to the effect that assertions of vague sentences can trigger accommodation of the range of application of vague predicates (Kyburg & Morreau (2000) and Shapiro (2003)). The question is whether accommodation happens in the way suggested by the Accommodation Strategy. There is some linguistic evidence suggesting that it doesn't.

It is quite natural to think that *if* we are dealing here with an instance of accommodation, it is an instance of presupposition accommodation (for other kinds of

[9]For a more technical definition of gappy contexts see Pagin (2010) and Pagin (2011). For the purposes of this paper we don't need to complicate our definitions that much.

[10]For some considerations that weight against accommodation see Gauker (2008).

accommodation see Lewis (1979))[11]. After all, according to this strategy, we accommodate by coming to presuppose something—that makes the context set gappy in the right way—in order to avert defectiveness. This is precisely the mark of presupposition accommodation; when an assertion is made that imposes requirements on the common ground that are not met, participants in the conversation tend to accommodate by presupposing what's required in order to avoid defectiveness.

Consider a standard example. 'Bruno stopped smoking' presupposes that Bruno used to smoke. This sentence, the story goes, cannot have a truth-value unless it is in the common ground that Bruno used to smoke. Thus, if a speaker asserts this sentence in a context where it isn't presupposed that Bruno used to smoke, participants in the conversation, cooperative as they are, will tend to accommodate by coming to presuppose that Bruno used to smoke in order to avoid defectiveness.

The question now is: what exactly is presupposed by an utterance of (FA)? We have been told that it is a presupposition that makes the context set gappy. This narrows down the candidates, but we are still far away from a definite answer. Now, to be fair, it is better to interpret this view as holding that vague sentences presuppose one of many propositions that can turn the context set into a gappy one—the alternative is too implausible. After all, on this view, vague sentences require the context set to be gappy, and there is a number of propositions that we could presuppose for this purpose. Whatever proposition is selected may vary from context to context depending on the objective of the conversation and the participants involved. Thus, rather than thinking that (FA) rigidly presupposes a specific proposition—like 'The King of France is wise' presupposed that there is a King of France—it is better to think that it presupposes one among many propositions, and that the proposition selected may vary from context to context.

Let's consider some presupposition candidates:

1. Fodor's apartment is either clearly near or clearly far away from the Met.

[11] Notice that standard of precision accommodation, and other kinds of accommodation, won't deliver a gappy context set.

2. Fodor's apartment isn't a borderline case of 'near the Met'.

3. It's not indeterminate whether Fodor's apartment is near the Met.

4. It's not unclear weather Fodor's apartment is near the Met.

This isn't an exhaustive list of all the purported candidates, but it is representative enough to cast serious doubts on the Accommodation Strategy. It may look like if we were to presuppose any of (1)-(4), the result would be a gappy context set relative to (FA). After all, one may think, if we update the common ground with any of these propositions the context set will only contain worlds where Fodor's apartment is either clearly close or clearly far away from the Metropolitan. Relative to this context set an assertion of (FA) would be in perfect harmony with principle (II).

The problem with this strategy is, however, that one cannot update the common ground with any of (1)-(4) without violating principle (II). A context set that is soritical relative to (FA) is also soritical relative to any of (1)-(4)—this is so because a sorites series for $\ulcorner F \urcorner$ is also a sorites series for $\ulcorner Clearly F \urcorner$, $\ulcorner Borderline F \urcorner$, $\ulcorner Indeterminately F \urcorner$, or what have you. Thus, from the perspective of the Acommodation Strategy, updating the context set with (FA) is as problematic as updating with any of (1)-(4). There are worlds in the context set such that it is indeterminate whether (1) is true in them.[12] The same can be said of the other three propositions. Vagueness in conversation is difficult to eradicate.

Perhaps the Accommodation Strategy needs precision to combat vagueness. If the proposition presupposed by (FA) is precise, then we can accommodate with it, get a gappy context set, and be in line with principle (II). It isn't altogether clear what these propositions could be, but here are a two representative candidates:

5. Fodor's apartment is either within 200m from the Met or it is at least 10km away from the Met.

[12] Rayo (2008) defends himself against this worry by invoking his *instrumentalism*. As I argue elsewhere **?**, an appeal to instrumentalism is quite helpful when compatibility with the phenomenon of vagueness is at issue. When it is, and appeal to instrumentalism idealizes vagueness away.

6. Fodor's apartment is either within 8 blocks from the Met or at least 20 blocks away from the Met.

These two options don't face the same problem as (1)-(4); presumably we could use them to update the context set into a gappy one without violation of principle (II).[13] However, this kind of strategy faces different problems.

If a sentence $\ulcorner S \urcorner$ presupposes that p, and an assertion of $\ulcorner S \urcorner$ is accepted, then it is infelicitous to follow by asserting p—however, it isn't infelicitous to first assert p and then assert $\ulcorner S \urcorner$. Consider the following examples:

7. Bruno used to smoke and he stopped smoking.

8. # Bruno stopped smoking and he used to smoke.

9. I have a sister and I have to pick her up from the airport.

10. # I have to pick up my sister from the airport and I have a sister.

The reason why (7) and (9) are felicitous and (8) and (10) aren't is familiar (Lewis (1979)). As we know, for 'Bruno stopped smoking' to be acceptable, we need to presuppose that he used to smoke. This is why it is perfectly natural to first say that Bruno used to smoke and *then* that he stopped smoking. (9) is felicitous for a similar reason. Now, the reason why (8) isn't felicitous is that we can only accept its first conjunct if we first accept the second, but then the second conjunct is simply redundant. Uttering (8) isn't that different from uttering the awkward sounding 'Bruno used to smoke and he stopped smoking and Bruno used to smoke'. (10) is infelicitous for analogous reasons.

Now, one may think that this test may support the Accommodation Strategy, rather than the other way around; after all, the following two sentences sound quite awful.

---

[13] There is a strict sense in which even (5) and (6) face the same problem as (1)-(4). Presumably there are worlds in the soritical context set where it is indeterminate whether Fodor's apartment is within 200m from the Met. This could be so because, as the literature on the Problem of the Many (Lewis (1993) and Weatherson (2003)) makes clear, there is some indeterminacy regarding the exact physical extension of material objects. I suggest to ignore this problem for the time being for the sake of exploring a different kind of problem this strategy faces.

11. # Fodor's apartment is near the Met and his apartment is either within 200m from the Met or it is at least 10km away from the Met.

12. # Fodor's apartment is near the Met and his apartment is either within 8 blocks from the Met or at least 20 blocks away from the Met.

Notice, however, that the reason why these two sentences sound bad is very different from the reason why (8) and (10) sound bad. (11) and (12) sound bad for the following reason. When one asserts a disjunction of the form either $p \vee q$ one implicates that one doesn't know whether p is true or whether q is true, otherwise one would be violating the Maxim of Quantity (Grice (1970)). This is why it isn't felicitous to first assert something that is incompatible with one of the disjuncts and then assert the disjunction. For example, an assertion of something of the form $\neg p \wedge (p \vee q)$ is infelicitous—unless one is in the logic classroom.

Now, (FA) is definitely incompatible with 'Fodor's apartment is at least 10km away from the Metropolitan'.[14] This is why (11) sounds bad; we cannot felicitously assert the first sentence and then imply that we don't know whether it is true that Fodor's apartment is at least 10km away from the Metropolitan. In order to mark a sharp distinction between (8) and (11)—one where the infelicity of (8), but not (11), is related to what is presupposed—let's consider the following sentences:

13. # Bruno stopped smoking and he used to smoke or he didn't use to smoke.

14. # Bruno stopped smoking and he used to smoke.

15. # Fodor's apartment is near the Met and his apartment is within 200m from the Met or at least 10km away from the Met.

16. Fodor's apartment is near the Met and his apartment is within 200m from the Met.

The last sentence isn't like the others. There are many contexts in which an utterance of (16) is felicitous *and* informative. Certainly it would have been more helpful for

[14] Of course, context makes clear that walking distance is at issue.

the bystander to have uttered (16) rather than simply (FA), whereas it's hard to think of a situation where it would be helpful to utter 'Bruno stop smoking and he used to smoke'. This is good evidence to the effect that an assertion of (FA) doesn't presuppose that the apartment is within 200m from the Metropolitan or at least 10km away from the Metropolitan.[15] A similar reasoning can be launched against (6) and similar alternatives.

Perhaps (16) is felicitous for a slightly different reason. If one likes the Accommodation Strategy one may think that participants to the conversation accommodate a assertion of (FA) in slightly different ways. After all, many different propositions can deliver a gappy context set—think of all the slight variations over (5) and (6) that can make the context set gappy. Some of the participants may accommodate with (5), others with (6), and others with something else. Then, the reason why (16) is felicitous is because it serves the purpose of letting the audience know how one wishes them to accommodate. This could be useful because it would be a way of preventing a defective context set—a case where different participants in the conversation have different presuppositions.

In the abstract this sounds like a very reasonable explanation. It is when we consider the details that its plausibility becomes dubious. In order to entertain the following objection it is enough to consider the speaker's perspective. Imagine our helpful bystander right before she utters (FA). She is in a soritical context relative to this sentence. The Accommodation Strategy predicts that the speaker will come to presuppose something that will make the context set gappy. There are many things that she—and the other participants in the conversation—could presuppose. Let's suppose that she presupposes (5). It is likely, the story goes, that other conversational participants, after the utterance, will come to presuppose something slightly different from (5)—we are good at coordinating with each other, but not at this level of precision. If this is so, unless the speaker utters (16) the context set will be defective. That is why, according to this line of thought, asserting (16) is felicitous and informative; it is a way of averting

[15]Notice that if p presupposes $\neg p \vee q$, then p presupposes q.

defectiveness by providing information about what is being presupposed.

But if this is so, then it is a mystery why we don't utter sentences like (16) all the time. After all, cooperative speakers typically try to avoid defective contexts sets for the sake of avoiding misscommunication. Moreover, if the Accommodation Strategist is right, it is a mystery why we utter sentences like (FA) to begin with rather than sentences like this:

17. Fodor's apartment is within 200m from the Met.

If what the speaker wants is to update the context set in such a way that (17) is entailed by it, it is much more effective and clear to utter (17) right off the bat than it is to utter (FA). This is so to the point that uttering (FA) over (17)—or some other precise sentence—would constitute a violation of the maxim of clarity and quantity. Thus, if the Accommodation Strategy is right, it is a mystery why we don't utter sentences like (17) instead of sentences like (FA).

4.4 The Dynamic Strategy

Pragmatics alone cannot neutralize the effects of vagueness in discourse, or at least it can't do it in a good number of interesting cases.[16] We need a theory of linguistic communication that makes room for the phenomenon of vagueness in a direct way, without pragmatic evasions. Before exploring the Quietist way, it will be instructive to consider a different strategy. A dynamic theory in the spirit of Barker (2002) has many attractive features. Most notably, it offers a systematic account of metalinguistic uses of vague sentences; it explains why certain vague assertions provide information about discourse, rather than information about the world.[17] This theory is also well equipped to account for the phenomenon of faultless disagreement (Barker (2013)) and the paradox of clarity (Barker (2009)). Finally, it aims at making room for the

[16] However, as Kyburg & Morreau (2000), Manor (2006), Rayo (2008), and Shapiro (2006) show, pragmatics can carry a good amount of weight.

[17] For example, in certain contexts an assertion of, say, 'Bruno is tall' can convey new information about standards of absolute tallness and no new information about Bruno's height.

phenomenon of vague linguistic communication while being in line with principle (II).[18]

In this section I wish to challenge the claim that this theory does in fact make room for the phenomenon of vague linguistic communication. If my criticism is correct, (II) gets little support from this camp. First I shall briefly sketch the features of the theory that are of interest to our discussion. Then I shall proceed to argue against certain features of it. This discussion also serves the purpose of motivating the Quietist approach.

4.4.1 The Dynamics of Vagueness

Barker (2002) exploits the seemingly trivial fact that participants in the conversation presuppose that they are having a conversation (Stalnaker (1978), Stalnaker (1998)). Hence, every world in the context set is such that the relevant conversation unfolds .[19] This is a powerful tool, simply because there are many discourse features that we could exploit in theorizing about vagueness and other semantic phenomena. Let's take a look now at which features of discourse this theory exploits.

Given that the focus is on vague predicates, the feature of discourse this theory exploits is directly related to gradable adjectives. Following Lewis (1970), let d be a function that maps worlds onto the delination that gradable adjectives have in the discourse of that world. A deliniation is a function from adjectives to degrees. Thus, where w is a world, $d(w)([\![tall]\!])$ is the standard of absolute tallness in w. Technicalities aside, the thought is that 'tall' gets associated with a degree relative to a world, where this degree represent the standard of absolute tallness at that world. Lets assume that in w such a degree is 1.80m. It follows that everyone who is at least 1.80m is tall in w,

[18] Barker doesn't explicitly discuss principle (II), but his theory certainly assumes it.

[19] Thus, rather than modeling the context set as a set of worlds, we could do it as a set of pairs $\langle w, s\rangle$, such that w is a world and s is a discourse *in* that world. This way of representing the context set is a matter of choice. As Stalnaker (1998) pointed out, given that a discourse is always part of a world, every single aspect of a given discourse is already captured by the hosting world. As such, we don't *need* to explicitly include in the representation of the context set a discourse, assignment functions, discourse referents, or what have you. Thus, as before, we could very well model the context set as a set of possible worlds. We shall do it this way following Barker (2002)—however, notice that Barker (2009) and Barker (2013) uses pairs, rather than single worlds.

and the rest are not tall in w.[20]

The way this theory attempts to capture the vagueness of 'tall' and other gradable adjectives is by letting the deliniation function have different values at different worlds in the context set. For instance, 'tall' may have different standards of absolute tallness in different worlds in the context set. If we presuppose, as it is reasonable, that anyone who is at least 1.90m is tall, then the standard of absolute tallness will be lower than 1.90m in every world in the context set. If we presuppose that anyone who is shorter than 1.50m is not tall, then the standard of absolute tallness is greater than 1.50m in every world in the context set. Now, let's assume that it is common ground that Bruno is a solid 1.78m tall—so he has that height in ever world in the context set. Furthermore, assume that relative to our conversation he counts as borderline tall. The way in which this theory represents this is by modeling the context set as having some worlds where the standard of absolute tallness is grater than 1.78m and others where it is lower. As such, the context set doesn't entail that Bruno is tall, but it doesn't entail that he isn't tall either—this is so because there are some worlds in the context set where he is tall and others where he isn't tall.

Suppose now that you want to flatter Bruno, so you say 'Bruno is tall'.[21] If accepted, the effect of your assertion is, according to this theory, the elimination from the context set of all the worlds such that Bruno doesn't exceed the standard of absolute tallness of those worlds. For example, given that Bruno is 1.78m, the worlds where the standard of absolute tallness is at 1.79m have been excluded by your assertion, but the worlds where the standard is 1.77m are preserved. Now we are in a position to see how this theory is compatible with principle (II). Assertions of vague sentences have as an effect a determinate partition of the context set. This is made possible by the architecture

---

[20] Standards of comparison don't have an effect on the following arguments. As such, they will be ignored in order to keep the explanation simple.

[21] Barker's dynamic semantics interpretation of 'Bruno is tall' is this: $[\![\lambda C.\{w \in C : c \in \mathbf{tall}(\mathbf{d}(w)([\![tall]\!]), \mathbf{bruno})]\!]$, where $\mathbf{d}(w)([\![tall]\!])$ is the standard of absolute tallness in w and $\mathbf{tall}(\mathbf{d}(w)([\![tall]\!]), \mathbf{bruno})$ is the set of worlds w where Bruno is at least as tall as the standard of absolute tallness in w. Thus, the content of 'Bruno is tall' is a function that takes a context and delivers the set of worlds in that context where Bruno is at least as stall as the standard of absolute tallness. For the sake of simplicity I shall avoid using the dynamic interpretation of sentences; the static counterpart will do. Our discussion won't be affected by this decision.

of the theory; in each world in the context set either Bruno exceeds the standard of maximum tallness of that world or he doesn't.

It *may* appear as if this theory has no trouble accounting for Fodor's Apartment while holding to (II). Here is how the explanation goes. Relative to each world in the context set, the semantics assigns a degree to 'near the Metropolitan'. If the degree in w is, let's say, 200m, then anything within a 200m radius from the Met is near the Met in w. The relevant degree isn't the same in every possible world in the context set—this is intended to capture the fact that some locations are borderline cases of 'near the Met'. The distance between the apartment and the Met also varies across worlds in the context set. This captures the fact that we don't know—nor presuppose—the location of the apartment. Given this, the crucial thing to note is that every world in the context set is such that Fodor's apartment is either near the Metropolitan or not near the Metropolitan. Hence, an assertion of (FA), if accepted, has the effect of excluding all the worlds such that the apartment isn't within the relevant radius in those worlds. For example, it excludes all the worlds where the apartment is 202m away from the Metropolitan and the degree is lower (i.e., 200m), and it preserves all the worlds where the apartment is 202m away from the Metropolitan and the degree is higher (i.e., 203m).[22] It seems, then, that this theory delivers the right prediction; relative to our example, an assertion of (FA) is perfectly felicitous. What is surprising is that this theory delivers the right prediction without fancy pragmatic footwork and without sacrificing (II).

4.4.2 Dynamic Problems

Once the dust settles, the surprise goes away: vagueness had to be sacrificed. There are two places where one may feel that this theory has idealized vagueness away.[23] The first is in the assignment of precise degrees to gradable adjectives relative to worlds.

[22]Of course, it also excludes all the worlds where the apartment is 195m away from the Met and the degree is 193m, and it doesn't exclude the worlds where the apartment is 195m away and the degree is 196m.

[23]One can idealize vagueness away if one's theoretical purpose is, say, to explain certain metalinguistic effects of vague assertions. As such, the following criticism doesn't put into question core features of this theory.

The second is in the assumption that the the context set is always crisp and fully determinate—there is no indeterminacy regarding which worlds are in the context set. Let's consider each objection in turn.

Precise delineation functions are suspicious. What exactly does it mean to say that at w the maximal degree of tallness is 1.80m? If w is in the context set, then it is a live option—it is a way the world might be. But is it? If epistemicism is true, then yes, that is a way the world might be—it could even be the actual world for all we know. In fact, the easiest way to interpret this theory is from an epistemicist standpoint.[24] Alas, that is not how the theory is meant to be interpreted (Barker (2002)). But how can we interpret it? If vagueness is a real semantic phenomenon there is nothing in w that could determine a sharp standard of absolute tallness. Hence, such a standard cannot be determined by w. But if so, how can w be excluded from the context set solely on the basis of the standard of absolute tallness associated with it?[25] This theory explains indeterminacy as variability of standards across worlds in the context set, but it does that at the cost of making each of the worlds very hard to understand. In a sense this theory grasps the Supervaluationist's precisifications and promotes them to live possibilities. Whatever the technical benefits of such a move, its philosophical justification is unclear.

Now, the assumption regarding standards of tallness holds quite a bit of weight. This theory has an account of (FA) that is compatible with (II) in good part because it assumes that the delineation function assigns a specific standard of absolute tallness to 'tall'—and other gradable adjectives—relative to each world in the context set. It is in good part because of this assumption that in each world in the context set Fodor's apartment is either near the Metropolitan or not near the Metropolitan. It seems that this theory can preserve (II) at the cost of assuming that, at a world, it isn't vague

[24] The idea being that in the actual world there is a precise standard of absolute tallness, and the context set having worlds with different standards represents the fact that we don't know—nor presuppose—what the exact standard is. This, for sure, would be a natural way of being an epistemicist and a stalnakerian.

[25] Recall that after you assert 'Bruno is tall' worlds where the standard of absolute tallness is greater than 1.79m get excluded. So these worlds get excluded form the context set because of the standard they have.

what the standard of tallness is. Unless we can make sense of this assumption in non-epistemicist terms, we should question this explanation's plausibility.[26]

There is a different problem. This one runs much deeper, given that it springs from the theory's architecture. As it stands this account assumes that the context set is crisp and perfectly determinate. There seems to be no indeterminacy as to whether some worlds are in the context set, and therefore no indeterminacy regarding what the live options are. However, this kind of indeterminacy is abundant when we use vague languages. Preparing a common ground for conversation is a messy affair; it is similar to collectively packing a joint suitcase under time pressure before storming out to the airport. However, this theory makes it looks as if the construction of the common ground is the result of meticulous work and negotiation.

In our run-of-the-mill conversations we may presuppose, for instance, that Magic Johnson is tall (for a person, not a basketball player). And then, perhaps, we also presuppose that people *a bit shorter* than him are tall as well. This is the material we use to build the common ground, and its vague. If this is what we presuppose—along with other vague things—we should expect that our common ground entails that Magic Jonson is tall, and also that someone a millimeter shorter than him is tall as well. And the same goes for other people a bit shorter than this. However, it cannot be the case that this context set entails that someone is tall, but that it fails to entail that someone who is a millimeter shorter is tall. Such is the nature of vagueness.

However, things in the dynamic framework are not this way. This is so because this theory assumes that the delineation function has a determinate range of variability. By this I simply mean that, on this account, every context set C is such that there is precise degree n such that 'tall' doesn't have a standard of absolute tallness greater than it in

[26] Notice that supervaluating over the delineation function won't help the cause. Suppose that instead of assigning a single delineation function at a world we assign a range of admissible ones. So if a person is tall at w relative to all the admissible delineation functions, then she is tall. But if she is only tall relative to some of them and not others, then it is indeterminate whether she is tall. Going this way helps alleviate the objection we are entertaining, but it does so at the cost of sacrificing (II). This is so because on this revised version of the theory an assertion of (FA) is indeterminate in some worlds in the context set—those worlds where the apartment is within the relevant distance relative to some of the admissible delineations but not others. Of course, on my view (II) has to be sacrificed, but for now we are considering theories that attempt to preserve this principle.

any $w \in C$. Similarly, there is a degree n' such that 'tall' doesn't have a standard lower than it in any $w \in C$. Hence, for each context set there is a precise pair of numbers n and n' such that the range of degree variability is between them—the standards of absolute tallness at a world in C is somewhere between n and n'. For example, a context set C can be such that the degree variability range for 'tall' is between 1.80 and 1.50. Thus, this context set entails that anyone who is 1.80m or taller is tall, but it doesn't entail that someone who is 1.7999m is tall. This seems to be a false consequence of the theory, given that we usually construct our common grounds by presupposing vague things, i.e., Magic Johnson is tall and people a bit shorter than him are tall as well.

Given these features, the dynamic theory has some bad predictions regarding informativeness. For example, any conversation with context set C is such that it *is* presupposed that anyone who is at least 1.80m is tall, but it *isn't* presupposed that anyone who is 1.799m is tall.[27] It is hard to believe that if this conversation is conducted using vague terms, then an assertion of 'anyone who is at least 1.80m tall is tall' would be completely uninformative, whereas an assertion of 'anyone who is 1.7999m tall is tall' would be informative—the first assertion couldn't be used to eliminate worlds from this context set, but the second could. Of course, there could be a conversation like that, but it would require a very specific set up that our run-of-the-mill conversations lack. If there is semantic vagueness regarding the standard of tallness—and there is—this theory doesn't make enough room for it.

4.5 Other Alternatives

Before exploring the Quietist alternative it would be instructive to consider a few other options. The first one we shall consider is similar in some respects to the Accommodation Strategy and similar in other respects to the Dynamic view. After considering this view we shall discuss whether the Stalnakerian framework could benefit from standard theories of vagueness, like Supervaluationism and Degree theory.

---

[27]It could be argued that this is just an idealization that shouldn't be taken seriously. I could agree with this, so long as vagueness isn't at issue. Given that vagueness is at issue, one should make this kinds of idealizations.

4.5.1 The Diagonalization Strategy

One may think that we could treat vague assertions in a way similar to how Stalnaker (1978) treats the Zsa Zsa Gabor example. (In the next few paragraphs I shall assume that the reader is familiar with Stalnaker's diagonalization strategy.)[28] The idea is that vague sentences express bivalent propositions relative to a context, and that these proposition vary from context to context. In one context a vague sentence containing 'tall' can express a proposition with a certain standard of tallness, and in other contexts the same sentence can express a different proposition with a different standard of tallness—these sentences draw a cut-off at different locations in different contexts.

Now, the thought goes, if vague sentences are context dependent in this way, we (almost) never know—nor presuppose—what context we are in. As such, the context set should include worlds containing each of these contexts. But then, after an assertion of 'Bruno is tall' this sentence doesn't express the *same* proposition in each world in the context set—in worlds with one kind of context the sentence expresses one proposition and in worlds with another kind of context the same sentence expresses a different one. This assertion would violate principle (III), unless we diagonalize. The result of diagonalizing is the preservation of worlds w where Bruno meets the standard of tallness determined by the context in w, and the exclusion of worlds w' where Bruno doesn't meet the standard of tallness determined by the context in w'. Hence, the thought goes, there is an incentive to diagonalize; doing so guarantees that the conversation is in line with principles (I), (II), and (III). This view is similar to the Accommodation Strategy in that both views appeal to pragmatic mechanisms in order to preserve principle (II).

This view has a flavor similar to Barker's approach. Both theories predict that the same worlds get excluded as a result of asserting 'Bruno is tall'—however, they arrive at this from different angles.[29] As expected, these views also suffer from similar problems. When discussing Barker's view we where perplexed by the fact that, according to it,

[28] Thanks to Dirk Kindermann (p.c) for suggesting this idea.

[29] The main difference (abstracting away from the dynamic/static differences) is that on Barker's view the same function from context to context gets expressed in every world, whereas in the view we are discussing a sentence expresses different propositions in different worlds. This requires that one view requires diagonalization whereas the other doesn't.

delineation functions assign a precise standard of absolute tallness relative to worlds in the context set. Now we can be perplexed at the diagonalization view assigning precise bivalent propositions relative to a context. This could be easily understood from an epistemicist perspective, but it is difficult to make sense of if we hold that vagueness is a semantic phenomenon. This view also assumes that it is determinate what worlds are in the context set. But, once again, if there is vagueness concerning standards of tallness, it should be indeterminate whether some worlds are in the context set. This strategy is in line with principle (II), but it can only do so at the cost of distorting the phenomenon of vagueness.

4.5.2 Standard Theories of Vagueness

Principle (II) is hopeless. We cannot both preserve it and make justice to the fact that conversations are usually conducted using vague languages.[30] Hence, we have to consider some ways of making the Stalnakerian framework vagueness friendly without worrying about principle (II).

A natural thought is to incorporate a standard theory of vagueness into this framework. To this end, we could use insights from Three-Valued Theory (Tappenden (1993) and Tye (1994)), Supervaluationism (Fine (1975) and McGee & McLaughlin (1995)), Degree Theory (Machina (1976) and Edgington (1997)), or Contextualism (Kamp (1981a), Gaifman (2002), Soames (1999), Raffman (1994), and Shapiro (2006)).[31] For instance, we could supervaluate over the context set—thereby claiming that there are multiple admissible context sets and that it is indeterminate which one is the right one. This could help alleviate *some* of the problems. For instance, this way of doing things recognizes that some worlds are such that it is indeterminate whether they are in the context set: the worlds that are in some, but not all, of the admissible context sets. Or perhaps, following Degree Theory insights, we could claim that membership in the context set is a matter of degree. This could also help alleviate *some* of the problems

[30]This, of course, under the assumption that vagueness a semantic—rather than an epistemic—phenomenon.

[31]Mention Gómez-Torrente (2010), Manor (2006), and Unger (1979) as a different kind of alternative.

concerning indeterminate membership to the context set. These are just two examples among many.

All these theories have valuable contributions to our understanding of the phenomenon of vagueness. However, as I have argued elsewhere **?**, all of them suffer from an important problem; given that they assume Exhaustiveness, they fail to make room for the phenomenon of vagueness.

- **Exhaustiveness**: For each object o in the domain and for each predicate $\ulcorner$F$\urcorner$ in the language, the semantic status of o with respect to $\ulcorner$F$\urcorner$ is fully specified; either $\ulcorner$F$\urcorner$ applies to o, or it doesn't, or this is, to some degree, indeterminate.

Any theory that abides by this principle cannot make justice to the fact that vague predicates do not draw sharp cut-offs relative to soritical domains. A full argument for this claim is quite complex, but it easy to convey the core ideas supporting it.

Picture a sorites series for the predicate 'is blue'. Let's say that the first tile in the series is bright blue and the last one is bright purple. A view that assumes Exhaustiveness would try to specify, for each member of the series, what its semantic status is with respect to this predicate. Now, independently of the details of such a view, there has to be a first member of the series that is a assigned a semantic status that is less than clearly/determinately/absolutely/indubitably blue. We can be confident that vague predicates do not draw this kind of classification—this is, in fact, their distinctive characteristic. Modeling a predicate as drawing this kind of cut-offs is to represent it as not being vague; at best a predicate like this can be indeterminate, if it has borderline cases. This is so whether we take the vague predicate to be drawing a vague classification relative to a series of tiles or possible worlds; recall that a context set can be soritical relative to a sentences containing a vague predicate. This is a major theoretical cost, *if* one of the objectives of our model is to be compatible with the phenomenon of vagueness. As such, it isn't altogether clear that we should incorporate one of this theories into the Stalnakerian framework of linguistic communication—these theories may not be what we need to make sense of the idea that the context set is often vague.

Defenders of standard views have tried to alleviate this problem by appealing to the phenomenon of higher-order vagueness. There are excellent reasons to to be dissatisfied with those attempts (Chapter 2, Sainsbury (1996), Zardini (2013),Tye (1994), Fine (2008), Fara (2003), Gómez-Torrente (2002)). The problem, in a nutshell, is that these theories either fall pray to higher-order vagueness paradox, or they have to postulate sharp cut-offs at high-orders of vagueness. There is no plausible way of justifying any of these alternatives. Given this, we should be very suspicious about standard theories of vagueness, in so far as they assume Exhaustiveness. If Exhaustiveness is the root of the problem, we need to rely on a theory of vagueness that doesn't hold to this principle. *Quietism* is an excellent alternative to consider.

4.6 Semantic Quietism

According to Quietism it is in principle impossible to specify the full range of application of a vague predicate relative to a soritical domain. This aspect of the view constitutes a sharp denial of Exhaustiveness and, as we shall see, it is the first step towards securing compatibility with the phenomenon of vagueness. On this view, only a partial specification of the range of application of vague predicates can be given. Now, and this is important, the Quietist thinks that the proper way to be partial is to be *silent*; there are some members of a soritical domain such that our semantic models have to be completely silent as to whether or not the relevant vague predicate applies to them, or whether this is, to some degree, indeterminate.

The Quietist trades a picture where the range of application of a vague predicate is understood in terms of positive, borderline, and negative regions, for one where there are positive, silence, and negative regions. Without introducing crucial subtleties it doesn't appear as if the Quietist framework achieves any progress. A Quietist model has a sharp division between the positive region and the silence region just as a standard theory has a sharp division between the positive region and the borderline region. One may think that if standard theories ought to be rejected for that reason, then one should reject Quietism as well. This is, however, too quick—some subtleties can make a significant difference.

When a Quietist model is silent as to whether a predicate applies to an object, this is compatible with the predicate in fact applying to the object. It is also compatible with the predicate not applying that object or with this being indeterminate. Hence, if a Quietist model is silent as to whether 'tall' applies to Bruno, this is compatible with Bruno being tall, Bruno not being tall, or Bruno being borderline tall. This is the first subtle difference between Quietism and standard theories. The second one concerns the point at which a Quietist model becomes silent. A model that secures compatibility with the phenomenon of vagueness has to be silent about certain semantic facts. For instance, there has to be someone who is, say, tall and that the model is silent about it. Similarly, there has to be someone who is not tall and that the model is silent about it. If this is done right, we can guarantee that our semantic model is compatible with the phenomenon of vagueness.

The way to achieve this is by having a silence region between the positive and the negative cases, and to make sure that this silence region contains some positive cases and some negative cases. Here is an example of this (where the vague predicate at issue is 'blond' and the red lines mark the silence range):

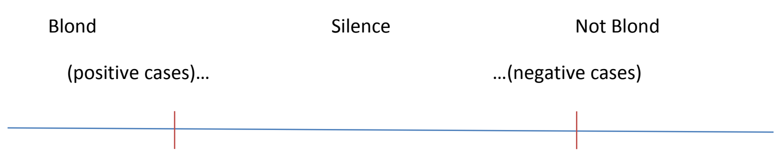

This critical move guarantees that the model's cut-offs land between either two positive cases or two negative cases. As such, these cut-offs are inoffensive—they only represent the point at which the theorist decided to become silent. Standard theories of vagueness misrepresent the vague transition from positive to negative cases by drawing sharp cut-offs where there are none. Quietism makes room for the phenomenon of vagueness by becoming silent before being forced to misrepresent vague transition from positive to negative cases.

A good Quietst model, then, is one that is silent about certain semantic facts.

Thus, in a sense, a good Quietist model has to be *sub-optimal*. This is so because any Quietist model M that secures compatibility with the phenomenon of vagueness is such that there is another model just like it that can capture at least one more semantic fact. For example, if Olivia is tall and M is silent about it, there can be a model M' that is just like M except because it specifies that 'tall' applies to Olivia. If so, then, M' is better than M, in so far as it captures at least one more semantic fact. There is no way around it, a Quietist knows that any model that she may want to offer is such that there is another one better than it. This is the price we pay for compatibility with the phenomenon of vagueness.

This raises the question: how to select a Quietist model? For sure there are many sub-optimal models, so is there one that we ought to chose? Is there such a thing as *the* best Quietist model? A positive answer to these questions would imply that there is a precise place along a sorites series where good models ought to become silent. This, however, is absurd. The Quietist thinks that the concept of *the best Quietist model* is not a useful one. Rather, we should focus on Quietist models that are *good enough* for our theoretical purposes. There are many models that are good enough given certain purposes. Some of them are better than others. Any of them will do, so long as all of them are good enough. Of course, from a Quietist perspective, the project of finding *all* the models that are good enough is as misguided as the project of finding *the* best model. We should rest content with identifying some models that are good enough—however, all we need is one.

In what follows we shall focus on how a Stalnakerian theory of linguistic communication can benefit from Quietist insights. One interesting result is that the project of modeling a given conversation puts constrains on what counts as a good enough Quietist model. This makes it the case that selecting a Quietist model isn't a completely arbitrary affair. Towards the end we shall reflect on how to live without principle (II).

4.7 Quietism and Conversations

Given that Quietism is the only theory out there that can guarantee compatibility with the phenomenon of vagueness, it is a good idea understand how this theory can be incorporated into the Stalnakerian framework. If these two theories can be integrated, the result is an account of linguistic communication that is compatible with the phenomenon of vagueness. As we have seen, no other theory on the market can promise this much. But, as it is evident at first glance, this theory also comes at a price—a Quietistic model is, by its own nature, sub-optimal. However, as we shall see, reflective equilibrium favors the Quietist way.

Let's get into some of the details. Recall that the context set gets determined by standing presuppositions. For the most part, presuppositions are stated using vague vocabulary—we can presuppose that you like *red* cars, that my sister is *tall*, that Fodor's house is *near* the Metropolitan, that someone you know is *rich*, and so on. Getting a non-vague context set out of these presuppositions would require heavy-duty pragmatic work during conversation—however, we never work that hard during our run-off-the-mill conversations. It shouldn't come as a surprise if the context set exhibits vagueness. Similarly, the sentences we assert during conversation are usually vague. As such, we should expect their contents—or the relationship between those sentences and several precise contents—to be vague.[32] Given this, the first order of buisness is to understand how to model the content of vague sentences in a Quietist way. After doing so, the representation of the context set will be straightforward.

In a sense, declarative sentences can be used to draw *classifications* in modal space. The sentence 'Olivia is tall' can be used to draw a classification separating the worlds where this sentence is true from those where it isn't. If the sentence is vague, we should expect this classification to be vague as well. It's a good thing that now we know how to model vague classifications in a Quietist framework. We can represent the content of this sentence as a classification that has been incompletely specified. Thus, for example, we can represent it as a classification including all the worlds where Olivia is 1.80m or

[32]Say why I want to remain neutral about this.

taller, as excluding all the worlds where she is 1.50m or shorter, and as being silent about the rest. Of course, the worlds where Olivia is 1.799m are sauch that she is tall, however, this specification is silent about it—this is just one among many incomplete specifications. What is important is that this Quietist representation confines to the silence range the intractable transition from the positive to the negative cases. As expected, this Quietist representation of this content is compatible with the fact that the set of worlds where this sentence is true is vague. As we know, compatibility comes at a price; we have to be silent about certain robust semantic facts. This isn't a surprise, given that, according to Quietism, it is in principle impossible to capture every significant aspect of a conversation conducted using vague terms.[33] This is another instance of a Quietist rejection of Exhaustiveness.

In a Stalnakerian framework the context set is the intersection of all the presuppositions in play—recall that presuppositions are being modeled as sets of possible worlds. Quietism can agree on this, so long as we accept that vague presuppositions can only be incompletely specified. As such, the context set has to be incompletely specified as well. Let's say that our model is silent as to whether someone who is between 1.76m is tall and whether someone who makes 260k dollars is rich. Furthermore, let's assume that we presuppose that my sister is *tall* and that that you know someone *rich*. If this is so, according to Quietism, our model has to be silent as to whether the worlds where my sister is 1.76m and the worlds where the person you know makes 260k are in the context set. The outcome is a context set that has only been only partially described. If the theorist chooses her silence region wisely, she would have a description of the context set that is robust enough to predict the evolution of the conversation, and silent enough to make room for the phenomenon of vagueness. This isn't an ideal position, but we can't do better than this.

[33]Something to note is that one need not subscribe to Quietism's metaphysical claims and still be able to model vagueness in a Quietist way. Thus, one may doubt that it is in principle impossible to model the range of application of vague predicates in an exhaustive way and still have good reasons to use the Quietist insights. Thanks to Agustín Rayo for pointing this out to me.

Of course, we need a rigorous understanding of how the intersection operation between two sets that have been incompletely specified works—otherwise we cannot understand assertion in a Quietist framework. Appendix 1 explains how to intersect sets of possible worlds that have only been incompletely specified. Setting those technicalities aside, it suffices to say that the result of intersecting incomplete specifications is, more often than not, another incomplete specification—silence breeds more silence.

The Quietist treatment of Fodor's Apartment is straightforward. Let's say that our Quietist model is silent as to whether something that is anywhere between 200m and 400m from the Metropolitan is *near* the Metropolitan.[34] Now, previous to the assertion of (FA) the context set is soritical with respect to this sentence. After the participants in the conversation accept this assertion the Quietist models the context set as excluding worlds where the apartment is 400m from the Metropolitan or further away, and remains silent with respect to worlds where the apartment is anywhere between 200m and 400m away from the Metropolitan. The rest of the worlds are, of course, in the context set. The Quietist can only model the evolution of this conversation in an incomplete way, and by doing so in the right way she opens up enough room for the phenomenon of vagueness.

One may also wonder what is, according to Quietism, the logic of natural languages. Quietism is fairly neutral with respect to this question. What logic we chose depends in good part on what we want to say about the Sorites Paradox and penumbral connections. In Chapter 3 I show how certain Supervaluationist insights can be incorporated into the Quietist framework. The result is a view I call *Superquietism*. This view preserves classical logic and validates penumbral connections for the fragment of the language that is D free, while being compatible with the phenomenon of vagueness. As such, on this view 'Bruno is tall or not tall' is a tautology and 'if Bruno is tall, then he isn't short' is always true. Superquietism isn't the only way to go. However, other alternatives remain unexplored.

There are may pressing questions that remain open. In what follows we shall focus

[34] Of course, if this is a good Quietist model something that is 200m away from the Metropolitan *is* near the Metropolitan.

on some details that need to be provided. In particular the two questions that need to be answered is how to select one Quietist model among many and how assertions can be felicitous even if they fault principle (II).

4.7.1 When to Become Silent?

We can specify some aspects of the context set, but we also need to be silent about others. How should we decide what to be silent about? Our answer to this question has the same flavor as our discussion in the previous section; there is nothing like *the* best way to represent the context set. In particular, for each conversation, there is no unique number of worlds such that the model has to be silent about them. Thus, we can only offer a *sub-optimal* model of a given conversation—assuming that the conversation in question is as wild and vague as our everyday conversations. This is so because every good Quietistic representation of the context set is silent about some worlds that are in fact in the context set. This, of course, doesn't mean that we have no guidance when it comes to deciding which Quietistic representation of the context set to use. We want to be silent enough to make room for vagueness, but we also want to be informative enough to predict the evolution of the context set. This criterion—vaguely—narrows down the number of Quietistic representations that can be useful.

Let's consider a simple example that illustrates some of the issues that may come up when modeling conversations in a Quietistic way. Let's say we are getting to know each other and we've been having a conversation for a little while. So far I have informed you that I'm from Mexico City, that I have a sister, that I'm a chess enthusiast, and that my dog's name is 'Owen'. Having seen how tall I am, you ask whether my sister is tall. I reply by saying: 'Yes, Ale is tall (for a female)'. The Quietist faces some decision points; what I have asserted is vague, so the Quietist model has to be silent about whether or not certain worlds have been excluded from the context set. But which worlds? There is, of course, no unique answer to this question. But there are certain considerations that can help us narrow down our options.

For instance, it would be a bad idea to let the model be silent about worlds where my sister has the height she actually has. Doing so would be to represent purely my communicative intentions—I have said, truly, that she is tall. My sister's height is 1.75m, so it would be good practice to let the Quietist model be such that all the worlds where she is 1.75m and taller are worlds where she is tall. This criterion helps us identify some points where the model *shouldn't* be silent. Now, if someone who is 1.75m, is tall, then someone who is 1.7499 is also tall. It may be a good idea not to be silent about those worlds—they are extremely similar to worlds that we shouldn't be silent about. By this criterion we may want to be specific, rather than silent, about worlds where she is 1.74m and, perhaps, even worlds where she is 1.73m (she is tall in all those worlds). After taking into account considerations of this sort it is hard to find strong reasons for the model not to become silent about the world where Ale is, let's say, 1.7256m.[35] The Quietist modeling our conversation could very well decide to be silent about worlds where Ale is between 1.7256m and 1.60m; doing so would be semi-arbitrary, but that is all we can aspire to in this position. She could also decide to become silent about worlds where Ale is 1.7299m and 1.61m; this would also be a semi-arbitrary decision. There are, of course, many other semi-arbitrary decisions of this kind that could be taken.

What matters is that the Quietist model is compatible with the phenomenon of vagueness *and* that it captures the communicative intentions of the speakers along with other central aspects of linguistic communication. Once the Quietist has secured this, she may feel free to make semi-arbitrary decisions of the kind discussed. These decision are harmless, so long as we realize that at some point the theorist *may* have to take them back in order to accommodate future assertions. For example, assume that the theorist modeling our conversation decided to be silent about worlds where Ale is between 1.7256m and 1.60m. Now, suppose that at a later stage in our conversation Andrea, a friend of yours, joins us. What you say by way of introduction is: "Andrea is 1.71m. I enjoy having tall friends like her." A semi-arbitrary decision was made, and it has to be taken back in order to accommodate your assertion. The theorist decided to be silent about whether worlds where a female is 1.71 are worlds where she

[35] Of course, in richer conversations where we have talked extensively about people's height, this may be different. However, this is just a difference in complexity, not in kind.

is tall. In order to capture your communicative intentions the theorist cannot remain silent about this. Thus, she has to adjust her model, and make other semi-arbitrary decisions—perhaps now she may want to be silent about whether worlds where a female is between 1.68m and 1.60m are worlds where she is tall.[36]

4.7.2 Life Without Principle (II)

We have gathered good reasons to think that principle (II) is untenable. We have also presented a picture of vague linguistic communication that conflicts with it. It is time to confront the motivations for this principle; if we are to reject it, we need to understand what is wrong with the ideas supporting it.

According to Stalnaker (1978), the point of an assertion is to reduce the context set in a *determinate* way. If the speaker fails to do so—by asserting a sentence that is vague relative to the context set—she fails to make her intentions clear. If so, the story goes, her audience won't know how to update the context set, given that the sentence asserted is indeterminate relative to some worlds in the context set. This will induce a conversational crash or some sort of pragmatic repair. If this is correct, it is a good idea to accept principle (II). Where does the mistake lie?

We can all agree that the point of an assertion of a precise sentence is to reduce the context set in a determinate way. But why think that the point of a vague assertion is to do the exact same thing? There is a point in asserting vague sentences; we do it all the time and we do it for good reasons. To think that the objective of a vague assertion is to update the context set in a determinate way is to misunderstand them all together.[37] It would appear that in a good number of cases we assert vague sentences *because* we are not in a position to update the context set in a determinate way—given

[36] This rises interesting questions regarding forced march situations. There is no unique way in which a Quietist can respond to a situation like this. My personal favorite is this: if someone asks you for each member of a sorites series whether that member is, let's say, tall, then the wise thing to do is to become silent before you run out of positive cases. It is important to recognize that there is no unique point where you must become silent. How far you go is just a mater of how risk averse you are. Now, if someone is holding a gun against your head, and she is ready to shoot if you make a mistake or become silent, then you won't live for long.

[37] It goes without saying that the main focus in Stalnaker's 'Assertion' wasn't on assertions of vague sentences. This is probably why principle (II) didn't seem problematic.

our conversational interests and standard practical limitations. Of course, there is some theoretical interest in working with an idealized framework where agents are perfectly rational and don't face practical limitations as we do. Nevertheless, this isn't—or at the very least it shouldn't be—the objective of a theory of vague linguistic communication. We use vague vocabulary, in good part, because usually our conversational goals are rough and we face considerable practical limitations.

Let's explore these ideas in more detail. It is clear that in a good number of cases we are not in a position to update the context set in a determinate way—or at the very least not without being uncooperative. However in a good number of those cases we may still be in a position to provide useful information. A good example of this is Fodor's Apartment. In this example the helpful bystander doesn't know exactly where the apartment is—this is a practical limitation. However she has a vague, and useful, piece of information; she knows that the apartment is *near* the Metropolitan. When she asserts (FA) her objective is not to update the context set in a determinate way; she is in no position to do such a thing. For sure, there is no determinate set of possibilities that she means to exclude with her assertion. Her objective is to be as useful and informative as she can given the position she is in.[38]

There are other cases where, even if we are in a position to update the context set in a determinate way, it is better not to do it. Let's say that our conversational goal is to get to know each other. As usual, we have limited time to do so—this is a practical limitation. Suppose you are curious as to whether my family members are tall. I'm in a position to be very precise; I could tell you exactly how tall each of them is, and then I could answer your question while making explicit each of the relevant standards of comparison—tall for an adult male in central Mexico and tall for an adult female in central Mexico. Or I could simply say 'yes, they are tall'. The precise answer overburdens the common ground with unwanted information. As a cooperative conversationalist I don't want to do that. The vague assertion is the most efficient way to answer your question in an informative way. This is why even if I'm in a position

[38]She could make a random guess and say that the apartment is on the corner of such-and-such, but then she would be quite uncooperative.

be precise—and make the context set evolve in a determinate way—it is better to be vague. There is a point in asserting vague sentences, and it isn't to update the context set in a determinate way.[39]

Another claim that is used to motivate principle (II) is that if the sentence asserted is indeterminate relative to some worlds in the context set the intentions of the speaker won't be clear. If so, the story goes, the audience won't know how to update the context; it won't be clear whether those worlds should be excluded from the context set. Now, we can all agree that if the sentence asserted is indeterminate in this way, the intentions of the speaker won't be *perfectly* clear. However, given the practical limitations discussed above, we shouldn't expect speakers, not even perfectly rational ones, to have perfectly clear linguistic intentions. As we have seen, very often it's better to assert something vague. Hence, very often it's better to have less than perfectly clear linguistic intentions; very often it *shouldn't* be clear whether we intended to exclude certain possibilities. Principles of rational communication should be concerned with whether speakers are clear enough, and not on whether they are perfectly clear. The helpful bystander in Fodor's Apartment was clear enough (given her practical limitations). When in our conversation I told you 'yes, they are tall', I was clear enough as well.

We have been focusing on the speaker's perspective, but we should also reflect on the audience. If the intentions of the speaker are not perfectly clear—if her assertion is indeterminate relative to some live options—how can the audience update the context set? This is another concern motivating principle (II).[40] Now, this would be a problem if, as it were, the audience had to eliminate worlds *manually.* If they had to hold worlds in their hands and decide whether to knock them out, less than perfectly clear speaker intentions would be problematic. But in a Stalnakerian framework we eliminate worlds from the context set by updating the common ground. And we do this simply by acquiring new presuppositions. Once the participants in the conversation come to presuppose, say, that Fodor's apartment is near the Metropolitan, the particular way

[39] Rayo (2008) offers examples where speakers update the context set in a determinate by way of using vague vocabulary. Although I agree with his assessment, those examples are the exception, rather than the norm.

[40] Thanks to Carlotta Pavese for pressing this point.

in which the context set gets updated is vaguely determined by what they presuppose.

Perhaps what is problematic isn't the update, but the audience's perspective on the context set. If the intentions of the speaker are unclear, and the audience updates with the relevant assertion, it won't be clear to them which possibilities are still live options. But this doesn't seem to be a problem at all. On the contrary, this is how things should be, so long as there are enough possibilities that are clearly live options. In Fodor's Apartment, after we update the common ground with (FA), we many not be clear about whether worlds where the apartment is 12 blocks away from the Metropolitan are still in the context set. However, so long as it is clear that, say, worlds where the apartment is withing 8 blocks from the Metropolitan are still live options—and perhaps some other worlds too—our perspective on which possibilities are still live options is detailed enough to carry on with our inquiry. Maybe as we march on we realize that we need to gather more information. If so, we may ask another helpful bystander for directions. But maybe we won't need further aid and we can find Fodor's apartment along way, where it should be, near the Metropolitan.

Appendix 1: Intersection, Union, and Difference

Intersecting sets of possible worlds—or anything else, for that matter—that have only been incompletely specified can be a bit laborious, but the strategy is quite intuitive. If Γ is a set of possible worlds that has been incompletely specified, then we can use a triple to represent it. The first member, P_Γ contains some of the worlds in Γ, S_Γ contains the worlds the Quietist model is silent about with respect to membership in Γ, and N_Γ contains some of the worlds that are not in Γ. If A and B are two sets that have been incompletely specified, we define intersections, union, and difference as follows (where "$\cup^q$", "$\cap^q$" and "$-^q$" are our Quietist operators and "$\cup$", "$\cap$", and "$-$" are borrowed from standard set theory):

- $A = \langle P_A, S_A, N_A \rangle$
- $B = \langle P_B, S_B, N_B \rangle$
- $A \cap^q B = \langle P_A \cap P_B, S_A \cup S_B - N_A \cup N_B, N_A \cup N_B \rangle$

- $A \cup^q B = \langle P_A \cup P_B, (S_A \cup S_B) - (P_A \cup P_B \cup N_A \cup N_B), N_A \cap N_B \rangle$
- $A -^q B = \langle P_A - P_B, S_A \cap S_B, (N_A \cup N_B) \cup (S_A \cup S_B - S_A \cap S_B) \rangle$

Something to note is that, from a formal perspective, we can do without S_Γ. We could simply stipulate that P_Γ and N_Γ are not mutually exhaustive.[41] The only reason I explicitly include S_Γ is because it is easier to interpret the formalism this way.

[41] Thanks to Agustín Rayo for pointing this out.

Bibliography

Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, *25*(1), 1–36.

Barker, C. (2009). Clarity and the grammar of skepticism. *Mind & Language*, *24*(3), 253–273.

Barker, C. (2013). Negotiating taste. *Inquiry*, *56*(2-3), 240–257.

Beaver, D. & Zeevat, H. (2004). Accommodation.

Cook, R. T. (2002). Vagueness and mathematical precision. *Mind*, *111*(442), 225–248.

Dorr, C. (2003). Vagueness without ignorance. *Philosophical Perspectives*, *17*(1), 83–113.

Dorr, C. (2010). Iterating definiteness. *Cuts and Clouds*, *1*(9), 550–586.

Edgington, D. (1997). Vagueness by degrees. In R. Keefe & P. Smith (Eds.), *Vagueness: A Reader*. MIT.

Fara, D. G. (2003). Gap principles, penumbral consequence, and infinitely higher-order vagueness.

Fara, D. G. (2010). Truth in a region. *Vagueness and Language Use. Palgrave Macmillan, Oxford*.

Fara Graff, D. (2000). Shifting sands: an interest relative theory of vagueness. *Philosophical Topics*, *28*(1).

Fara Graff, D. (2002). An anti-epistemicist consequence of margin for error semantics for knowledge. *Philosophy and Phenomenological Research*, *64*(1), 127–142.

Field, H. (2003). No fact of the matter. *Australasian Journal of Philosophy*, *81*(4), 457–480.

Fine, K. (1975). Vagueness, truth and logic. *Synthese*, *30*(3), 265–300.

Fine, K. (2008). The impossibility of vagueness. *Philosophical Perspectives*, *22*(1), 111–136.

Frege, G. (1879). Begriffsschrift, a formula language, modeled upon that of arithmetic, for pure thought. *From Frege to Gödel: A source book in mathematical logic*, *1931*, 1–82.

Frege, G. (1893). *Grundgesetze der Arithmetik: begriffsschriftlich abgeleitet*, volume 1. H. Pohle.

Frege, G. (1988). *Grundlagen der Arithmetik: Studienausgabe mit dem Text der Centenarausgabe*. Meiner Verlag.

Gaifman, H. (2002). Vagueness, tolerance and contextual logic.

Gauker, C. (2008). Zero tolerance for pragmatics. *Synthese*, *165*(3), 359–371.

Gómez-Torrente, M. (2002). Vagueness and margin for error principles*. *Philosophy and Phenomenological Research*, *64*(1), 107–125.

Gómez-Torrente, M. (2010). The sorites, linguistic preconceptions, and the dual picture of vagueness. *Cuts and Clouds: Vaguenesss, Its Nature and Its Logic*, 228.

Grice, H. P. (1970). *Logic and conversation*. Harvard Univ.

Groenendijk, J. & Stokhof, M. (1991). Dynamic predicate logic. *Linguistics and philosophy*, *14*(1), 39–100.

Heck, R. (1997). The julius caesar objection.

Heim, I. (1983). File change semantics and the familiarity theory of definiteness.

Kamp, H. (1975). Two theories about adjectives. *Formal semantics of natural language*, 123–155.

Kamp, H. (1981a). The paradox of the heap. In *Aspects of philosophical logic* (pp. 225–277). Springer.

Kamp, H. (1981b). A theory of truth and semantic representation. *Formal semantics-the essential readings*, 189–222.

Kaplan, D. (1989). Demonstratives. In J. Almong, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan*. Oxford: Oxford University Press.

Keefe, R. (2000). *Theories of vagueness*. Cambridge University Press.

Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, *30*(1), 1–45.

Kripke, S. A. (1963). Semantical considerations on modal logic.

Kyburg, A. & Morreau, M. (2000). Fitting words: Vague language in context. *Linguistics and Philosophy*, *23*(6), 577–597.

Lewis, D. (1975). Languages and language. *Minnesota Studies in the Philosophy of Science*, 3–35.

Lewis, D. (1979). Scorekeeping in a language game. *Journal of philosophical logic*, *8*(1), 339–359.

Lewis, D. (1993). Many, but almost one. *Ontology, Causality and Mind*, 23–42.

Ludlow, P. (2014). *Living words: meaning underdetermination and the dynamic lexicon*. Oxford University Press.

Machina, K. F. (1976). Truth, belief, and vagueness. *Journal of Philosophical logic*, *5*(1), 47–78.

Manor, R. (2006). Solving the heap. *Synthese*, *153*(2), 171–186.

McGee, V. & McLaughlin, B. (1995). Distinctions without a difference. *The Southern Journal of Philosophy*, *33*(S1), 203–251.

Pagin, P. (2010). Vagueness and central gaps. *Cuts and Clouds: Vagueness, its Nature and its Logic, Palgrave Macmillan.*

Pagin, P. (2011). Vagueness and domain restriction. *s include: Anton Benz, Gerhard Jäger and Robert van Rooij (editors)*, 283.

Raffman, D. (1994). Vagueness without paradox. *The Philosophical Review*, 41–74.

Raffman, D. (2014). *Unruly words: a study of vague language*. Oxford University Press.

Rayo, A. (2008). Vague representation. *Mind*, *117*(466), 329–373.

Sainsbury, R. M. (1996). Concepts without boundaries.

Shapiro, S. (2003). Vagueness and conversation. *Liars and Heaps*, 39–72.

Shapiro, S. (2006). *Vagueness in context*. Oxford University Press.

Smith, N. J. (2008). *Vagueness and degrees of truth*. Oxford University Press.

Soames, S. (1999). *Understanding truth*. Oxford University Press.

Sorensen, R. (2001). *Vagueness and contradiction*. Oxford University Press.

Stalnaker, R. (1978). Assertion. *Syntax and Semantics*, *9*(4).

Stalnaker, R. (1998). On the representation of context. *Journal of Logic, Language and Information*, *7*(1), 3–19.

Tappenden, J. (1993). The liar and sorites paradoxes: toward a unified treatment. *The Journal of Philosophy*, 551–577.

Tye, M. (1994). Sorites paradoxes and the semantics of vagueness. *Philosophical Perspectives*, 189–206.

Unger, P. (1979). There are no ordinary things. *Synthese*, *41*(2), 117–154.

Varzi, A. C. (2007). Supervaluationism and its logics. *Mind*, *116*(463), 633–676.

Von Fintel, K. (2008). What is presuppositioin accommodation, again? *Philosophical perspectives*, *22*(1), 137–170.

Weatherson, B. (2003). Many many problems. *The Philosophical Quarterly*, *53*(213), 481–501.

Weatherson, B. (2005). True, truer, truest. *Philosophical Studies*, *123*(1-2), 47–70.

Weatherson, B. (2010). Vagueness as indeterminacy. *Cuts and Clouds: Vaguenesss, its Nature and its Logic*, 77–90.

Williamson, T. (1994). *Vagueness*. Routledge.

Williamson, T. (1999). On the structure of higher-order vagueness. *Mind*, *108*(429), 127–143.

Wright, C. (1976). Language-mastery and the sorites paradox. *Truth and Meaning*, 223–47.

Wright, C. (1987). Further reflections on the sorites paradox. *Philosophical Topics*, *15*(1), 227–290.

Wright, C. (2001). On being in a quandary. relativism vagueness logical revisionism. *Mind*, *110*(437), 45–98.

Wright, C. (2003). Vagueness: A fifth column approach. *Liars and Heaps*, 84–105.

Zardini, E. (2013). Higher-order sorites paradox. *Journal of philosophical logic*, *42*(1), 25–48.