Stand Over There, Please

The Dynamics of Vagueness, the Origins of Vagueness, and How Pie-Cutting Relates to Ancient Heaps of Sand

Seth Cable

Master of Science Thesis
Master of Science Programme in Logic
Institute for Logic, Language and Computation
University of Amsterdam
August 2002

Although many are thanked in the Preface,
this thesis is dedicated to *only* the following individuals:

Michael and Jane Cable, who, at some point
during the Carter administration, thought it wise to make a little boy.

Mary Cable, who con't spel, but has, concentrated at the tip of her pen, more
insight into a person's guts than I have in the whole of my fat, balding head.

Suzanne Kalamar, who submitted herself to heroic personal sacrifice, just so
she could spend a single harried week in a country without SUV's.
God in heaven, Zsuzsa, I hope you don't drop off the face of my Earth.

H. L. Mencken, whose timeless brilliance in the American idiom
and genuine intellectual bravery made existence tolerable
during those useless minutes before and after bed.

Contents

Preface

This thesis is an investigation into the semantics of vague expressions, although vague predicates are predominantly discussed. It proposes an analysis of vague predicates within the framework of dynamic semantics, and argues that this analysis predicts the observed relationship between the properties of having borderline cases, of being susceptible to sorites paradoxes, and of being context-dependent.

The first, introductory chapter concerns the questions of what, precisely, a "theory of vagueness" is and what it can hope to achieve. The material here is intended to provide the lens through which the more technical work of the following chapters should be viewed, and it is hoped that the characterization offered here of a program for the study of vagueness clarifies and unifies the existing literature on this subject. In outline, the chapter explains first the concept of "vagueness" as it appears in the linguistic and philosophical literature, and then discusses the necessary limits to any scientific study of this property. In regards to those limits, it is argued that a complete explanation of the semantics and logic of vague language is logically impossible, and that the most we as theorists can hope for is a set of increasingly refined approximations. These necessarily limitations, moreover, affect how one should describe the study of vagueness as an intellectual activity, and so how the reader should understand the material of the next three chapters.

In Chapter 2, the basic semantic analysis for vague predicates is proposed. This analysis uses the formal techniques of dynamic semantics to capture the unique information change potential of predications employing vague language. Moreover, the analysis is argued to correctly predict the relationship between the context dependency of a predicate and its having "borderline cases." Some basic properties of the resulting logic of vague predicates are then discussed, including its wide retention of classical logic and its ability to grab hold of "penumbral connections." Finally, the vagueness operator "definitely" is introduced, as well as its proposed information change potential, and some effects on our logic of adding the operator to our language are proven.

Chapter 3 begins with a host of difficulties for the basic analysis of the previous chapter. In particular, the analysis allows for the possibility of unnaturally precise uses of vague predicates and fails to predict the relationship between the presence of borderline cases and so-called "higher-order vagueness." It is suggested that these problems may be overcome if a more sophisticated theory of the origins of linguistic vagueness is adopted. The concept of a "vague selection" is introduced, as well as the hypothesis that all vagueness in language may be reduced to the "vagueness" of these selections. This hypothesis is spelled out formally within our dynamic semantics, and the resulting analysis is demonstrated to overcome the challenges raised at the chapter's beginning.

The fourth and final chapter concerns the sorites paradox. The basic questions concerning the sorites are reviewed, as well as some reasons against accepting it as a genuine paradox. It is then argued that the reductive theory of vagueness put forth in the previous chapter provides a unique perspective on what is occurring within the sorites argument, a perspective from which one spies a potential "resolution" to the paradox. Finally, it is argued that this analysis of the sorites, when combined with our reductive theory of vagueness, predicts that predicates are susceptible to sorites arguments if and only if they are vague.

Acknowledgements

Writing this thesis was, for scattered reasons, a great personal joy, my delight in these ideas perhaps revealing itself at times in the density of the text. Stoked by my acknowledged weakness for ostentation (alá Don King), the furnaces of my obsequiousness rattle from an overheated desire to wax lyrical over the many people and institutions which were midwife to me during this difficult labor [1]. Nevertheless, I keep it short.

Highest thanks go to Professor Frank Veltman, who supervised my research on this project. His expertise in these subjects supplied direction when it was lost and levied criticism where it was most needed. He also read my myriad of meandering project reports, which, as the reader may already imagine, could become annoyingly hectic. I wish here to also extend thanks to the other members of my defense committee, Professors Jeroen Groenendijk and Johan van Benthem.

Thanks are also due to the Fulbright Foundation, the Netherland-America Foundation, The Netherlands America Commission for Educational Exchange, the Institute for Logic, Language and Computation (h.f. "ILLC") and the Spinoza Project, all of whom at some point in the past year gave me lots of money.

The ILLC must also be singled out for individual thanks. The reasons for gratitude should be obvious, but let me state for the record that the ILLC is a singularly stimulating research environment, one in which I have grown and developed in ways I would not have before imagined, and which paid me the goodwill of letting me in its doors. The ILLC and the Universiteit van Amsterdam are the main characters of the heretofore most fascinating drama in my life, and they shall never be forgotten.

The material making up Chapter 2 of this thesis was presented at the 2002 Szklarska Poreba Workshop on the Roots of Pragmasemantics. Thanks go to

[1] Self-mockery is the highest form of apology.

the participants of the workshop, who offered valuable comments on a nascent form of that work. Special thanks also to Nick Asher, who raised an important point regarding the truth tables on pages 38 - 41.

Personal thanks go to the following unforgettable teachers, friends and colleagues (in backwards alphabetical order): Henk Zeevat, Domenico Zambella, Yde Venema, Raj Singh, Börkur Sigurbjornsson, Josh Sack, Will Rose, Ravi Rajani, Anna Pilatova, Berna Martinez, Irwin Lim, Juan Heguiabehere, Dick de Jongh, Willem Conradie, Luciano Buratto, Steffen Bauer, Roger Antonsen.

A final shoutout goes to my peeps back in Jersey.

Chapter 1

Introduction

What Kind of a Question is "What is Vagueness"?

1.1 The Meaning of "Vague"

One can rely on thoughtful people to balk at the mention of a "theory of vagueness." To the properly critical and educated, the very idea that semantic theory should accept "vagueness" as an object of study sounds as perverse as physiology treating of "ugliness" or "stupidity" becoming a subject of intense concern to cognitive psychologists. For all the world, it appears as if the prejudices of the investigator are tainting and threatening the objectivity of his conclusions.

When a linguist speaks of an expression being "vague," however, the evaluatory connotations of this description are dropped, and the word becomes instead a term of art. In the dialect of the linguist, the word "vague" denotes a natural class of expressions; the study of vagueness is simply the study of this class. Such study attempts to discover a common, underlying nature in the elements of the class, one which would explain the existence of an assortment of properties they often manifest.

This distinct term of art is, moreover, not so far removed from the common *koine*. Everyone seems to know approximately what is meant by describing a predicate, quantifier or adverb as "vague." This impression may be due to selective sampling amongst those educated in academic philosophy, but even the unspoilt laity would agree with linguists and philosophers in categorizing such words as "hot," "tall," "heap," and "adult" as vague, while withholding the description from "circle," "even" as it applies to numbers, and "kilowatt."

Despite this resemblance to its vulgar cousin, it is best to lay out the ground rules governing the label "vague" as it is used in the philosophic and linguistic circles. In these circles, and in this thesis, "vague" denotes a class of expressions which have three closely related properties: the presence of borderline cases, susceptibility to sorites paradoxes, and a standard-sensitive context dependency.

1.1.1 Borderline Cases

Perhaps the simplest and most vivid feature of vague expressions is that they admit of "borderline cases," entities for which it is necessarily uncertain whether the expression applies to them. C. S. Peirce wrote as early as 1902 that "a proposition is vague when there are possible states of things concerning which it is intrinsically uncertain whether, had they been contemplated by the speaker, he would have regarded them as excluded or allowed by the proposition."

It is a matter of some delicacy how the adverbials "necessarily" and "intrinsically" in the descriptions above are to be further unpacked. First, it should be noted that this necessity is always relative to a particular context; the same object may be a borderline case for a vague predicate in one context, but not in another. This point will be made clearer in Section 1.1.3. Assuming our context to be fixed, the most common explanation is that a borderline case is "necessarily" or "intrinsically" uncertain in the sense that its uncertainty is not a result of some lack of knowledge. For example, one might offer an SUV as a borderline case for the predicate "truck." Should an SUV be regarded as a truck? Speakers of English do not know. They have no strong intuitions on the matter, and one has a desire to say that this is despite their knowing all the relevant facts concerning SUV's: their size, their general body shape, what their most common uses are. Similarly, we may not be able to decide of a person exactly six feet in height whether he is "tall" or "not tall," and presumably no further information about the height of that person or the heights of others will allow us to make a decision.

Although it is a commonplace to regard the uncertainty characteristic of borderline cases as not arising from ignorance, certain authors have protested that to do so is tendentious and begs the question against theories of vagueness which claim borderline cases to be the result of an ignorance of the full meaning of vague expressions [1]. In the interests of fairness, then, the full force of "necessarily" and "intrinsically" in our definition above will for now be left unpacked. It will be assumed that the examples offered above communicate to the reader clearly enough what is intended by the notion of a "borderline case." Further elucidation of this notion may be gained from the following "behavioral criteria" offered by Graff 2000.

> "We are prompted to regard a thing as a borderline case of a predicate when it elicits in us one of a variety of related verbal behaviors. When asked, for example, whether a particular man is nice, we may give what can be called a *hedging* response. Hedging responses include: 'He's nice*ish*,' 'Well, it depends on how you look at it,' 'I wouldn't say he's nice, I wouldn't say he's not nice,' 'It could go either way,' 'He's kind of in between,' 'It's not that clear-cut,' and even, 'He's a borderline case.' If it is demanded that a 'yes' or 'no' response is required, we may feel that neither answer would be quite correct, that there is no 'fact of the matter'." (Graff 2000)

[1] Cargile 1969 and Williamson 1992 promote such accounts.

1.1.2 Susceptibility to Sorites Paradoxes

Although the presence of borderline cases is the simplest feature of vague expressions, sorites-susceptibility is both the oldest-known and the most widely-discussed. The sorites paradox derives its name from the ancient Greek word for "heap," historically the first vague expression to which the paradox was applied. All vague expressions, however, are subject to an argument of the following frustrating form: a single grain of sand does not constitute a heap; if a collection of sand is not a heap, then adding a single grain of sand to the collection does not thereby make it into a heap; therefore, any collection of sand is not a heap. Diogenes Laertius, for instance, demonstrates that an argument of this form may also be applied to the vague quantifier "few" [2].

> "It is not the case that two are few and three are not also. It is not the case that these are and four is not also (and so on up to ten thousand). But two are few: therefore ten thousand are also." (Barnes 1982)

Let us now more precisely state the claim that all vague expressions are "susceptible" to a sorites paradox. Although the following formulation covers only vague predicates, its generalization to other semantic types should not be difficult. If P is a predicate of language $\mathcal{L}$, then "P is susceptible to a sorites paradox" if there are entities α and β and a relation R such that the following sentences together strike a speaker of $\mathcal{L}$ as jointly satisfiable.

1. $P\alpha$
2. $\neg P\beta$
3. $\forall x \forall y((Px \wedge Rxy) \rightarrow Py)$
4. $\exists a_1 \ldots a_n(R\alpha a_1 \wedge Ra_1 a_2 \wedge Ra_2 a_3 \wedge \ldots \wedge Ra_n \beta)$

The claim that all vague predicates are sorites-susceptible, then, is the claim that for all vague predicates P in a language $\mathcal{L}$, one can for any speaker s of $\mathcal{L}$ find α, β, R which appear to s to be able to simultaneously satisfy conditions 1 - 4 above.

One can easily use this claim to identify vague predicates in one's own language. For example, that the vague predicate "tall" is susceptible to a sorites paradox may be seen by taking α to be Michael Jordan, β to be Michael Jackson, and R to be the relation "is one millimeter taller than". The "absolute positive" premise that Michael Jordan is tall and the "absolute negative" premise that Michael Jackson is not tall are both obviously true, and so jointly satisfiable. Moreover, the "inductive premise" that "if x is tall and x is one millimeter taller than y, then y is tall" also seems intuitively true. Finally, one is loathe to dispute the (often implicit) "existential premise" that one can imagine there

[2] The reader may observe that Diogenes here describes that form of the sorites in our century dubbed "Wang's Paradox" and famously discussed in Dummett 1975.

being a "chain" of individuals, each one millimeter taller than the other, linking Michael Jordan at one end with Michael Jackson at the other. This property is also easily confirmed for such predicates as "fat," "in New Jersey," "chair," and for any color term. For some predicates, such as "nice," stating precisely the identity of R becomes a challenge, but their sorites-susceptibility nevertheless appears upheld.

What makes the sorites-susceptibility of a predicate paradoxical is that, although a speaker finds for some α, β, R, sentences 1 - 4 to be manifestly consistent, these sentences are, by the manifestly sound principles of classical logic, jointly inconsistent. The speaker thus finds himself caught in an uncomfortable dilemma: either the manifestly sound principles of classical logic are unsound when operating on vague language, or in any imaginable situation one of the manifestly consistent 1 - 4 must be false. Either of these possibilities strikes one as frankly beyond the pale. This, of course, sets a tantalizing challenge before the philosophically minded, and it is fair to say that most literature on the subject of vagueness includes an attempt to find for the speaker some comfortable way out of the sorites dilemma. This thesis will not break from tradition, and in Chapter 4 such a "resolution" is proposed.

1.1.3 Standard-Sensitive Context Dependency

It has been observed of expressions having borderline cases or susceptible to sorites paradoxes that their meaning in a context depends upon the selection of a particular "standard." Although authors do not uniformly acknowledge the importance of this observation, nor place equal weight upon it, it nevertheless appears well-supported that the meanings of vague expressions vary with the context, and in an intuitive sense, they depend upon the "standards" the speakers are employing in the context for whatever is denoted by the expression.

To illustrate, consider that whether one describes a particular man as "tall" depends upon the standards one is employing: whether one is holding him up to the class of all men, all basketball players, or all Mediterranians. If one says "the Koninklijk Paleis is big," the truth of this statement depends upon whether, in the context, "big" means "big for a palace," "big for a structure," or even "big for a building located at Dam Square." Each of these standards for "bigness" is employed in a different context. Vague quantifiers such as "many" are also context-dependent in this way. If I were to say to you "Baskin Robbins has many flavors of ice cream," you may wonder whether I meant by my statement that Baskin Robbins serves more flavors of ice cream than is usual for a restaurant chain, that they serve more flavors of ice cream than one is accustomed to seeing anywhere, or that they serve more flavors of ice cream than what is necessary for our present purposes. Vague nouns seem not to offer counterexamples either. A large leather bag filled with beans is described as a "chair" when in the den, but certainly not when in the living room. All these cases demonstrate that the interpretation of a vague expression depends upon an implicit "comparison class," one which may change with the context.

The context dependency of vague expressions, however, extends beyond even this dependence upon a contextually determined comparison class. One may fix the comparison class for an expression and yet still find that its interpretation varies from situation to situation. There is a certain degree of arbitrariness or "whim" associated with the use of a vague expression. This is especially clear when one considers objects near the "borderline" for the expression. If one were to ask me, for example, whether Maxima Zorreguieta is "pretty for a woman her age," although the comparison class is explicitly fixed, my response will largely vary with the context: how I am feeling that morning, whether I want to be especially "picky," or even whether I have recently seen a particularly gorgeous woman all can influence my response. Thus, one may argue that it would be incorrect to identify the contextually varying standard for a vague expression with that expression's contextually given comparison class. Whether one agrees to this or not, the point remains that there is some notion of "standard" whereby it is true to say that the interpretation of any vague expression within a context depends upon a standard that varies with the context.

When linguists and philosophers speak of "vague expressions," then, they are speaking of a class of expressions which, along with some number of yet undiscovered properties, have borderline cases, are susceptible to sorites paradoxes, and whose interpretation in a context depends upon the selection of a "standard." This specialized use of the adjective "vague" diverges slightly from what one might find in everyday language. Keefe & Smith 1996a make the following clarificatory distinctions.

In informal conversation, "vague" is sometimes used to communicate a property linguists and philosophers prefer to name "underspecificity." The statement "Someone said something," for example, might sometimes be criticized as being "too vague." Similarly, "I'm thinking of some number larger than thirty" would usually be described as unhelpfully vague information about the contents of one's thoughts. As Keefe and Smith rightly point out, "vague" is in such cases being used to communicate underspecificity, the property of not being as informative as the context demands [3]. This property seems to have nothing to do with "vagueness" as understood by linguists and philosophers. "Some number larger than thirty," for example, does not admit of borderlines, is not susceptible to sorites arguments, and has a meaning which does not at all depend upon "standards" or change with the context. Likewise, the statement "Pauly lives in Manalapan, New Jersey," although vague in the technical sense, may in some context supply all the information one needs.

Vagueness should also be distinguished from ambiguity. Saying that someone's choice of words is "deliberately vague" sometimes means that that person has cleverly chosen to express themselves in ambiguous language: consider someone answering "yes" to the question "So, you didn't take the money?" An expression can be vague, however, without being ambiguous. The noun "toddler" is unquestionably vague, but it is also unquestionably univocal. Ex-

[3] In other words, "being vague" in this sense is simply violating the Gricean Maxim of Quantity.

pressions, moreover, may be ambiguous without being vague: a "lattice," for example, may or may not, depending on the definition one chooses to adopt, have a top and a bottom element.

These clarificatory distinctions, however, should not be taken as apodictic. That is, despite what one might infer from our making these distinctions, the properties of having borderlines, sorites-susceptibility and standard-sensitive context dependency should not be understood as anything like *criteria* for the application of the term "vague." "Vague," as a concept in linguistic theory, is not *defined* as the conjunction of these properties. Rather, it is believed that expressions with any one of these properties form a natural class, and the term "vague" is intended to denote *whatever* that class may be discovered to be. For the linguist, it merely happens to be a quite interesting fact that these properties appear to "hang together" perfectly. Thus, the claim that all vague predicates have these three properties is intended as a falsifiable hypothesis, though knowing what evidence could be universally recognized as disproving it may require our learning more about this class of expressions. The point of the clarificatory distinctions above, then, is to say that a linguist or philosopher would not, given the current state of his knowledge, characterize an expression as vague simply because it was underspecific or ambiguous.

That linguists and philosophers *do* take vague expressions to form a natural class may be supported by a few considerations. The foremost is, of course, the simple fact that the properties of having borderlines, sorites-susceptibility and standard-sensitive context dependency appear to be in a natural, indivorceable correspondence. Furthermore, as Keefe & Smith 1996a reports, theorists may be observed to debate what the proper criteria are for the term "vague" (p.14). The literature thus abounds with possibly conflicting definitions for "vague," and yet all researchers appear to agree on the relevant examples, intuitions and data. Moreover, the tone of these scholarly debates suggests that "vagueness" is implicitly understood as an objective property inhering in the expressions discussed, and that the arguing theorists are battling not so much over definitions of a term of art, but over characterizations of the underlying nature of that property.

Less direct evidence that "vague" denotes a natural class is that the class of vague expressions behaves as other natural classes in linguistics, in that speakers appear to have an implicit, natural knowledge of its contents incongruous with it being some artificial concoction thrown together by theorists. To illustrate, ordinary speakers, such as those in an undergraduate philosophy class, can upon being given only a handful of examples and a suggestive definition correctly identify all other expressions as "vague" or "not-vague" (i.e. "precise"). Since this suggestive definition, usually along the lines of "can have borderline cases," greatly underdetermines the class of predicates intended [4], one concludes that these speakers tacitly intuit a yet undiscovered common nature to these predicates. That is, this singular ability of unindoctrinated speakers suggests that

[4]The degree of this underdetermination may be gleaned from Keefe & Smith 1996a's discussion of the proposed definitions of "vagueness" (Keefe & Smith 1996a, p. 14).

in identifying an expression as vague, one acts not on the definition handed out in lecture, but on one's implicit linguistic knowledge. This knowledge independently carves out these complementing classes and floats recommendations to the speaker as to which predicates group with which others, though the speaker himself cannot consciously identify the source of those recommendations. This view accords well with the phenomenology of making these distinctions in practice. When asked whether some expression is "vague" in the technical sense, the theorist will find that in drawing his conclusions he does not so much employ linguistic tests or gedankenexperiments as he does a kind of flat introspection of the expression's meaning. Somehow, I know immediately that "ambrosia" is a vague predicate while "asbestos" is not. Having made such classifications, one then asks the questions *How are they made?*, *What important difference do we all seem to intuit between these predicates?* It is here that one turns to these so-called "definitions" to evaluate their accuracy.

1.2 The Question

Thus, for the linguist, "vagueness" is simply that unknown property of an expression which leads it to exhibit borderline cases, sorites-susceptibility and standard-sensitive context dependency, and so it may be considered a proper subject of objective, scientific study. The question *What is vagueness?* is the inquiry into the common bond between those expressions classified as "vague." Like all questions of this sort, one chips away at its answer by isolating properties which vague expressions have been observed to have and asking "Why would *these* expressions have *those* properties?" By hitting upon a common feature from which this array of observed properties would follow, one thereby makes hypotheses into the "nature of vagueness." Again, like all studies of this sort, this activity proceeds in a piece-meal fashion. The entire set of properties discovered to be peculiar to vague expressions is never faced at a single time. Indeed, as more knowledge is accumulated, more hypotheses evaluated, this set grows in its size. The theorist therefore selects proper subsets of this assortment of properties and makes proposals as to how such collections of features may have a common source.

In broad, sloppy stokes, this is the scientific investigation of vagueness. However, one may feel here that something is amiss. Such a description would make this subject appear non-exceptional. However, there is some circumstantial evidence that vagueness as a linguistic property *is* rather exceptional, and that its study has a status quite distinct from that of NPI licensing. One often hears of vagueness being described as a "problem" or, worse, a "philosophical problem." Articles and book-chapters sag with such titles as "The Problem of Vagueness" and "The Impenetrable Mystery of Ontic Vagueness." Even such deep-thinking members of the philosophical pantheon as Wittgenstein, Dummett and Fine seem to have devoted considerable energy to the advancement of this subject; since when has such a team assembled before the finer intricacies of Aktionsart? In some instances, this "problem of vagueness" is nothing more than the program sketched above; the "problem" is simply a question, the question of the nature of vagueness. In most instances, however, this problem is accurately described as a foundational problem for natural language semantics. The recognition of that problem dates back at least to Bertrand Russell's speech to the Jowett Society (Russell 1923), the first work of modern philosophy to be written on this subject. The problem, put plainly, is that the existence of vague predicates is difficult to render consistent with the basic doctrines and assumptions of standard semantic theory.

1.3 The Problem

The problem centers on extensions. Elementary semantic theory states that predicates, such as adjectives and nouns, have as their semantic value "extensions," which are usually informally described as that set of entities to which the predicate applies. This basic assumption, nestled in the bedrock of nearly all semantic theory, presupposes that the things to which any predicate applies form a "set," a mathematical object whose nature is completely specified by the ZFC axiom system. The ZFC axiom system, however, makes use of classical logic, and there is, of course, no reason internal to the discipline of set theory for it to do otherwise. It is then a trivial theorem of the standard theory of sets that for all objects x and all sets s either $x \in s$ or $x \notin s$. Now, if the objects to which the vague predicate "truck" applies form a set, then all objects are either an element of that set or they are not. Thus, the statement that all predicates have extensions entails that all objects are either accurately described by the predicate "truck" or they are not. But, this seems patently false. Consider again our example of the SUV. One is uncomfortable in saying that "truck" does not apply to SUV's, and yet one is also reluctant to agree to the statement that SUV's are trucks. Is there even a fact of the matter whether "truck" applies to SUV's or not? It certainly seems as if there isn't, and thus it certainly seems as if the standard assumption that all predicates have extensions cannot be correct. In brief, when the fact that most predicates in natural language are vague is no longer ignored, the most basic assumptions of standard semantic theory make obviously false predictions.

The "problem of vagueness," then, is that contrary to the opening sentence of standard semantic theory, not all predicates have extensions, not all predicates denote sets. Many scientific disciplines are faced with foundational problems such as this; the standard assumptions regarding rationality and information in most economic theories provide an interesting parallel. Whenever such problems are noticed, they immediately beg the question of why the theory seems to work so well, despite its clear deviance from fact. For linguists, the reflex response to this question has been that standard semantic theory, just as any scientific discipline, rests on an *approximation* of linguistic reality. The statement that all predicates have extensions is said to be merely a simplifying assumption, much as frictionless planes and perfectly homogeneous speaker communities are simplifying assumptions over certain domains of study.

In order for this to be true, however, it must be that the real semantic value of vague predicates has qualities sufficiently like a set that semantic theory does just as well to characterize it as a set when describing its behavior in combination with such other expressions as "all," "or" and "possibly." *Prima facie*, this is quite plausible, but more investigation into the semantics of vague predicates will be needed before the depth of the "problem of vagueness" can be properly sounded. Such investigation must address the equally puzzling question "if vague predicates do not denote sets, then what do they denote?" That is, if we would like to truly vindicate standard semantic theory as merely resting on a "simplifying assumption," then we must first converge upon a richer, more detailed description of the semantic value of vague predicates.

On this latter issue, the literature seems to skip between just a few well-rehearsed proposals. The most familiar ideas are already recognizable to (and in some cases reinvented by) undergraduate students: vague predicates denote "fuzzy sets"; vague predicates denote pairs of positive and negative extensions; vague predicates denote sets of model completions. Perhaps one can better evaluate these classic proposals, and thereby gain more insight into the question itself, through knowledge of the lexical semantics of particular vague predicates. Given a tentative theory of the denotation of vague predicates, one would want to determine whether for particular vague predicates there exist proposals regarding their lexical meaning which predict that their semantic value is one instance of the value our theory predicts all vague predicates to bear.

The lexical semantics of particular vague predicates becomes important from another perspective as well. One sees that a consideration of the question "What do vague predicates denote?" rests to some extent on an answer to the deeper, prior question "What are the laws, principles, generalizations governing the application of vague predicates to objects?" In framing hypotheses about the denotations of vague predicates, it helps to have some beginning data and intuitions concerning how particular vague predicates are applied. Having some set of doctrines, trivial as they may be, on the matter of how speakers decide whether a vague predicate is to be employed in the description of an object will give us a starting direction on how to address the puzzling data involving sorites arguments and borderline cases. As an illustrative example of what I have in mind, the claim that vague predicates denote fuzzy sets is inspired, in part, by the observation that speakers regard particular vague predicates as applying to objects "in differing degrees." Similarly, the claim that vague predicates denote sets of completions is inspired by the observation that such predicates may be made "more precise," their indeterminacy lessened through a reduction in their borderline cases, though their being made completely precise may be an impossible occurrence. Finally, even the classical theory that all predicates denote sets follows from the observation that within mathematics, any object is either accurately described by a predicate or it is not.

1.4 Our Inscrutable Knowledge of Vagueness

1.4.1 Lexical Semantics and the Study of Vagueness

Our hypothetical plan of action, then, is to gain some insight into the question "what do vague predicates denote" by investigating the lexical semantics of particular vague predicates, the implicit knowledge speakers have of how to use the predicate to describe and classify objects. Thus, we want to consider such questions as "What does a speaker recognize in an object such that they can decide whether or not it falls under the predicate?", "What kind of criteria are employed?", "How is the decision reached?" When we ask these questions of vague predicates, however, we find ourselves abused by another deep foundational difficulty. This problem is most likely familiar to all, but in different forms. In the literature, one can doubtlessly find many arguments along the following lines; I would like simply to offer my own formulation of the criticism.

To find a comfortable way into the problem, let us first reject two methods of answering our questions above. For certain scholars' purposes, an answer to the question "what must a speaker know in order to use the vague predicate 'tall'" may be provided by a Davidsonian truth theory for some relevant fragment of the language studied. That is, the interests of these theorists may be satisfied by any recursive system which derives the statement " 'Dave is tall' is true iff Dave is tall." Our interests, however, go a bit further. Certainly, our semantics must be able to predict the Davidsonian biconditional, but it must do more besides. In the worst case, all a Davidsonian truth theory states is that speakers use "tall" to describe tall things. Although this statement is not strictly circular, it does not address the question of how speakers decide what objects are tall, which is really the intent of our questions above. That appeal to a Davidsonian truth theory here would be a complete *non-sequitur* becomes clear upon recalling that such a theory is intended as an answer to the conceptual question "what must speakers of the language know" rather than the psychological question "what do the speakers know?"

Another sort of response to our questions above is equally peculiar, but due to the robust regularity of its appearance in the literature, it ought to be indicated how it falls short. Some scholars advocate the use of vague expressions in the metalanguage descriptions of the meanings of vague predicates. Again, when the purposes of these scholars are just those for which Davidsonian truth theories are designed, then no charges may be pressed [5]. However, one does occasionally find proposals along the lines of "the predicate 'tall' describes an object which has a height *significantly* greater than what might be expected." The work of Graff 2000 on sorites series and Barker 2002 on the expression "definitely" approach statements of this sort. The problem with these proposals is the work which is being done by the vague adverbial "significantly"; too often the vagueness of this adverbial is intended as the source of the vagueness of these expressions. When these proposals are meant as characterizations of the speaker's knowledge of vague language, through which one may understand how

[5] Our remarks here, therefore, do not challenge the position put forth in Sainsbury 1996.

vague language manifests the properties it does, then they straightforwardly commit the crime of circularity.

That they do succumb to circularity may not be immediately obvious, and so a few words should perhaps be said in support of this point. It will help to contrast an analysis of the kind criticized with one which is not circular in this manner. The "fuzzy set theory" analysis characterizes one's knowledge of vague predicates as a function assigning to objects real numbers from the interval [0,1]. From this characterization, a number of properties peculiar to vague predicates, and relating to adjectival comparison and the sorites paradox, may be argued to follow. In contrast, the putative "significantly" analysis derives these mysterious properties from the very vagueness of the word "significantly." The predicate "tall" is predicted to admit of borderline cases precisely because the word "significantly" does, but the word "significantly" admits of borderline cases just because it's vague. The same holds for all the other essential characteristics of vague predicates which this analysis "predicts." The analysis assumes the existence of, and so leaves unexplained, exactly that knowledge it was meant to draw out: how one applies vague predicates to objects such that their peculiar properties follow [6] [7]. This point cannot be phrased more effectively than in Kamp's "The Paradox of the Heap."

> "...the value of any semantic analysis depends on the adequacy of the logic which underlies the metalanguage in which the analysis is carried out...But what is the logic of a language containing vague concepts? ...to adopt a vague metalanguage would mean either that we put the cart before the horse, or else by that by defaulting on the theorist's obligation to make the logic of his language explicit, fail to attach the cart to the horse in any way whatsoever."
> (Kamp 1981)

Neither of these two proposals, then, will be acceptable for our purposes as characterizations of the knowledge speakers have of the meaning of vague expressions. The point of introducing them was both to elucidate what those purposes are and to segue into the following difficulty. If we pause for a second and reflect on the arguments just made against the "significantly" analysis, they

[6] In line with this point, we should make the following clarification. Graff 2000 seeks ultimately to reduce the vagueness of language to the metaphysically prior vagueness of a speaker's purposes. It should be noted that there is no circularity in *this* sort of approach, since the vagueness of language is not appealed to in the proposed analysis. That is, the "significantly" analysis cannot be said to explain the properties of vague expressions, since it does not explain the properties of the vague adverbial "significantly." A full reduction of the vagueness of language to the "vagueness" of some other object, such as a purpose, would thus not be circular in the way criticized since the supposed "vagueness" of this object would be a property entirely distinct from linguistic vagueness as described in 1.1. This point will have a special significance in Chapter 3.

[7] To balance this point, let us also here acknowledge that the "significantly" analysis of Graff 2000 does succeed in reducing all linguistic vagueness to the vagueness of a single adverbial. Moreover, there may be a semantic analysis of that adverbial which non-circularly predicts it – and so all other vague predicates – to exhibit its characteristic properties. However, this envisioned analysis would be a tad on the "disjunctive" side.

seem to generalize and force us into a rather disturbing conclusion regarding the viability of our project as presently conceived.

1.4.2 The Inscrutability of Vagueness

What, precisely, is meant by "a characterization of our semantic knowledge of a vague predicate"? Let us take as our paradigmatic example of such a characterization the claim " 'tall' describes an object if its height is greater than what might be expected of an object of its class." A "characterization of our semantic knowledge of the vague predicate P," then, seems to be some complex metalanguage predicate Φ such that " 'P' describes an object if Φ" is true, and from this fact it follows that "P" has a number of properties associated with vague predicates. Moreover, it is reasonable to conclude that such a characterization predicts "P" to have these properties simply because Φ has them within the metalanguage. For example, our paradigmatic characterization predicts that "tall" is context dependent simply because "an object of its class" is context dependent. The reader is reminded that since the goal is to characterize vagueness and not context dependency, this characterization does not commit the crime of the "significantly" analysis.

Now, thus far we have followed standard scientific modesty and supposed that our analyses could only be expected to predict a specially selected subset of those properties peculiar to vague predicates. Suppose, however, that we had the knowledge of gods. As deities, could we then obtain a *complete* characterization for some vague predicate? Could we ever have a full characterization of the speaker's knowledge of the vague predicate "tall," one from which *all* its properties as a vague predicate would follow? Well, following exactly the reasoning detailed above, such a characterization could not essentially employ vague language. The characterization could, of course, be expressed in vague language, but the presence of vague language in the characterization must not be necessary for it to make the predictions it does. If the vagueness of the language in the characterization is not being used circularly to explain the properties of "tall" as a vague predicate, then it should be possible to replace the vague language of the characterization with precise language and not threaten the success of the characterization. Thus, if there were a complete, non-circular characterization of a speaker's knowledge of a vague predicate, it should be possible for that characterization to be expressed entirely in precise language.

We seem now to be at the brink of a *reductio*, however. If one accepts that a complex predicate Φ composed entirely of precise expressions should thereby be precise, then this putative characterization simply cannot exist. If a complete characterization of a speaker's knowledge of a vague predicate were expressible in fully precise language, then one would have a precise predicate Φ which in our metalanguage displayed *all* those properties peculiar to vague predicates. Φ would have borderline cases, sorites-susceptibility, standard-sensitive context dependency, and any other properties one could think of as being characteristic of vague predicates. In short, the result is a precise predicate exhibiting the properties essential to vague predicates, and so therefore vague. Assuming that the set of vague predicates and precise ones is disjoint, one arrives at a contradiction. The conclusion is that there can be no complete, non-circular characterization of the semantics, and therefore the logic, of vague predicates.

The immediate consequence of this argument is that one cannot hope to fully characterize the knowledge a speaker has of the semantics of vague predicates, much less that knowledge which accounts for the properties essential to vagueness. Moreover, the argument applies equally well to any attempts at answering the question "What do vague predicates denote?" If there were an answer to this question which did not rely circularly on the vagueness of its language, then there would be a description of the semantic value of vague predicates phrased in entirely precise language. Such a description would have the semantic value peculiar to vague predicates, but would be precise—a blatant contradiction. Assuming that "circular analyses" are simply not genuine analyses and that incomplete answers are not genuine answers, one concludes that there is no hope of finding answers to the questions "what do speakers know about the meaning of 'tall'," "what do vague predicates denote," and "what is vagueness?" Any answer would simply not be coherent, and so one begins to doubt whether even the questions themselves are.

Again, it is not claimed that this conclusion is an original insight. One finds in the literature a general consensus that the logic and the nature of vagueness cannot be explained in precise terms. The argument above is merely meant as another way of reminding ourselves why this conclusion is true. However, although the reader may agree with the conclusion, he may yet be skeptical of the reasoning employed. To make this clearer, the following is offered as a more general argument of the same form. If the sets S and S' are disjoint and closed under the operation $\circ$, and if the relation R is such that $\alpha R\beta$ implies $(\alpha \in S \equiv \beta \in S) \wedge (\alpha \in S' \equiv \beta \in S')$, then for no $\alpha \in S$ can $\alpha R\beta_1 \circ \ldots \circ \beta_n$ for $\beta_1 \ldots \beta_n \in S'$. If $\alpha R\beta_1 \circ \ldots \circ \beta_n$, then since S' is closed under $\circ$, $\alpha \in S'$ and so S and S' are not disjoint. Where S is the set of vague predicates, S' is the set of precise predicates, R is the relation "completely characterizes," and $\circ$ is shorthand for the usual operations of semantic composition, the argument above is instantiated. Of course, a mere elucidation of the logic behind our argument is not going to evaporate all skepticism, and so we will now address two worries which might be raised against our reasoning.

1.4.3 Objections to the Argument and Replies

The first objection concerns a possible unsavory consequence of the argument for studies of ontic vagueness within the subject of metaphysics. These studies often attempt to build a theory of "vague objects" or of vagueness as a property of "the world," and employ mathematical, precise language in doing so [8]. The worry is that our argument likewise demonstrates that theories which employ mathematical, precise language in the metaphysics of ontic vagueness should be impossible.

[8] Parsons and Woodruff 1995 is a good example of such work.

In response to this objection, we should first of all make one very important point entirely clear: our argument does not at all rule out precise language being used successfully in *partial* characterizations of the semantics of vague expressions. This point will be further elaborated in Section 1.5, but let us here simply point out that the *reductio* rests on the putative characterization being *complete*. Thus, for an arbitrarily large proper subset S of the set of properties essential to vague expressions, there may well be non-circular characterizations of the semantics of vague expressions using only precise language which predict that those expressions manifest the properties in S. Similarly, in the worst case, all our argument implies for the study of ontic vagueness is that there can be no *complete* analysis of "vague" objects using mathematical concepts. But, what analysis thus far proposed claims to be complete?

Moreover, this "worst case scenario" is thankfully not the case. The argument has no consequences for the study of ontic vagueness unless some possibly false assumptions are made regarding the relationship between vague language and "vague" objects. To begin, note that the purpose of these repeated scare quotes around "vague" as it applies to objects is to remind us not equivocate between "vagueness" as it is introduced in Section 1.1 and whatever is meant by "vagueness" in the context of metaphysics. Unless one wants to make sense of the idea of objects being subject to sorites paradoxes, linguistic vagueness and ontic vagueness must be regarded as quite distinct properties [9]. Bearing this distinction in mind, one finds that for the study of ontic vagueness an argument parallel to that weighing against the study of linguistic vagueness does not go through. The argument at one point requires the premise that a non-circular analysis of a vague expression cannot essentially employ vague language. There is no reason, however, to suppose that non-circular analyses of ontic vagueness must not essentially employ vague language. Unlike our "characterizations of semantic knowledge," our explanations of the properties of "vague" objects presumably would *not* work because the *language* in which they were expressed had those properties. In fact, a direct application of our argument to the study of ontic vagueness belies an equivocation between the "vagueness" of objects and linguistic vagueness. It is only if one adds the assumptions that precise language can only describe "precise" objects and that "vague" and "precise" objects constitute disjoint classes, that an argument similar to the one we propose may go through for the study of "vague" objects.

This is all to say that the limitations revealed by our proffered argument really are unique to the activity of semantic theorizing. On the usual description of this activity, one theorizes about the meaning of an expression by creating a "model" of its meaning within a metalanguage. Creating such a "model," however, requires that the metalanguage expressions describing it take on exactly those semantic properties which one is seeking to explain in the target expression. The following diagram attempts to compactly describe this unique state of affairs.

[9] This is not to say, of course, that a deep and necessary relationship does not exist between these properties.

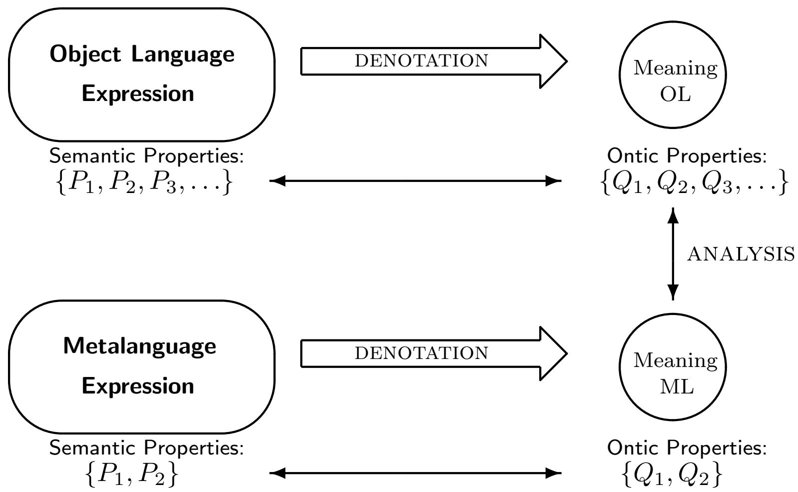

A second objection one may raise challenges the description of semantic analysis summarized in the above diagram and so crucial to the success of our argument. One may rightly be skeptical of the claim that in semantic theory the success of our analyses requires that they take on the properties of those expressions we hope to explain. After all, the greatest success story in the philosophy of language during the twentieth century was perhaps possible worlds semantics, a framework which permits semantic analyses of intensional language using only extensional machinery.

Contrary to the press releases of our discipline, however, possible worlds semantics still falls vastly short of a complete analysis of intensional language and so is simply grist for our mill. Recall that the disputed model of semantic analysis, when combined with our argument, entails only that *complete* characterizations for the targeted expressions are incoherent. The logic of intensional language, then, may be *partially* elucidated through the use of such non-intensional expressions as "set," "function," "relation" and so on. Quite important and otherwise baffling properties of the intensional expressions "possible," "believes" and "is allowed" may indeed be explained by characterizing the meanings of these expressions in non-intensional language. However, our argument predicts that there will *always* be properties of intensional language which will escape the extensional analysis. Thus, the well-known problems surrounding possible-worlds analyses of such intensional notions as "belief" and "necessary truth" will not dissipate. The semantic analysis of intensionality, however, need not be considered an absurd practice; Section 1.5 will describe one way of defending its coherency.

Let us now offer as a final point in favor of our argument that it predicts a particular sort of pattern in the attempts to answer questions concerning the nature of vagueness. Whatever hypotheses one's theory of vagueness casts regarding the semantic value of vague predicates, our argument shows that it can never predict those predicates to have the full range of properties they do. The most successful theory imaginable is one which predicts the finite subset of properties which theorists have so far observed. However, as soon as this theory is introduced, new properties should suddenly become apparent. The history of the study of vagueness, and in particular the "No Sharp Boundaries" paradox, follows exactly this grain. Once a new theory of vague language is published, there is observed a number of properties vague predicates have which are missing from the proposed model. Thus, logicians first proposed that all predicates, including the vague ones, denote sets. However, it was then quickly observed that vague predicates have borderline cases while set-denoting predicates cannot. Thus, supervaluational and many-valued logics introduced extension gaps into the semantics of vague predicates, in order to model the presence of borderline cases. However, then the peculiarities of higher-order vagueness were encountered, as it was observed that even the predicate "is a borderline case of P" admits of borderline cases. Thus, Fine proposed in 1975 an infinite hierarchy of supervaluational models, intended to reflect the infinite hierarchy of vagueness orders. However, it was then discovered that the very predicate "is an n^{th}-order borderline case of P for some $n < \omega$" was a vague predicate, admitting of borderline cases. This so-called "No Sharp Boundaries" paradox is, perhaps, one of a number of cases in the past and to yet come of the precise language necessary for a non-circular analysis of vagueness chasing the properties of vague predicates endlessly on.

1.5 Consequences for Our Study

There can be no answer to the question "What is vagueness?" There can be no genuine theory ascribing a property to vague predicates from which all the properties peculiar to them would follow. What, then, is left of the study of vagueness? To gain proper perspective, let us take a step back and recall what motivated this study in the first place.

Vague predicates display a number of properties which mark them as distinctly different from precise predicates. Many of these properties appear to "hang together" remarkably well. The study of vagueness begins from a curiosity over this co-occurrence of properties; we would like to know why *these* properties go with *those*.

Now, we cannot hope to understand the conjunction of all these relationships, but we can hope to understand, in a limited sense, a restricted subset of them. Our investigation of this restricted set of relationships, however, is in a unique way conceptual and not empirical. This point may, perhaps, be illustrated through the following interpretation of the philosophical literature on vagueness. The supervaluational analysis (Fine 1975), for example, helps us to understand how a predicate admitting of borderline cases, and thus violating bivalence, *could* nevertheless be subject to the Law of Excluded middle. Although this theory may be rightly criticized as leaving obscure the necessary relationship between first-order vagueness and higher-order vagueness, it addresses the mystery of penumbral connections in a way which singles it out over the fuzzy logic analysis. We may furthermore hope to improve upon the supervaluational analysis, just as we might an analysis in any empirical field, by offering an analysis which can elucidate all the relationships covered by it as well as the problem of higher order vagueness. This new analysis is guaranteed to be incomplete, and what we gain is only an understanding of *how* there *could* exist these various relationships, not *why* they *do* exist.

More concretely, I propose the following as a more accurate description of the linguistic study of vagueness than what was first presented in Section 1.2. Let $V = \{P_1, P_2, \ldots\}$ be the set of properties essential to vague language. Of course, only a very small proper subset of V is known to us now. Although we cannot understand how a predicate of natural language may have all the properties in V, we may select subsets of V and ask the more conceptual, philosophical question "how *could* a predicate come to have this set of properties?" To all outside appearance, our discipline carries on just as would a discipline concerned with truth, with questions of fact having objective answers. Just as in any of those disciplines, our discipline will always be troubled by a number of phenomena which we simply cannot yet account for. Just as in any of those disciplines, our discipline will seek to expand its knowledge by proposing hypotheses concerning its object of study which account for increasingly greater portions of its properties. Those hypotheses, again, will be rejected if discovered to be circular, explaining the properties of vague predicates by appealing to the vague language in which they are stated.

However, unlike those disciplines correctly labeled "scientific," the goal of our study of vagueness is not the resolution of a matter of fact. Although our study appears at any particular time to be organized as a scientific one, the limit point of all scientific disciplines is some entirely accurate description of their object of study, one through which all theoretical questions about its properties could be settled. For the reasons already outlined, our study cannot have such a limit point, and for this reason its short-term goals must also be differently defined. We cannot in all seriousness claim that our discussion is somehow converging towards a fully accurate characterization of a speaker's knowledge of vague language. Such a characterization simply does not exist. Rather, we should understand our discipline as "removing the mystery" of vague language. At any given time, there will be a proper subset of the properties of vague predicates which will strike us as simply mystifying. For example, given our present knowledge, it may still strike us as just perplexing that predicates with standard-sensitive context dependency *must* thereby also have borderline cases. We can "improve" our "understanding," relieve our worry, by removing this mystery. We may propose model languages, as linguists always do, in which such puzzling combinations of properties may be understood as issuing from a single source. We remove our dumbfoundedness, our complete frustration and awe, but at every step we ought not delude ourselves that there is some complete story regarding vagueness we are carefully and gradually unraveling. Given that the goal of this project is, in a certain sense, to console rather than to obtain some objective knowledge, it would perhaps be best described as a "therapeutic" activity.

And yet, our inquiry has all the external appearance of a scientific study: its use of problem solving, its standards for theory comparison and evaluation, perhaps even its outward semblance of progress. Seeing as how it will appear at any stage of its lifetime as a typical, scientific, linguistic study, I propose to continue on just as if it were. In the following chapters, arguments will proceed precisely as if there were some full, correct answer to the question "What is vagueness?" In particular, a theory of the nature of vagueness will be proposed which will appear to explain the relationship between borderline cases, sorites-susceptibility and standard-sensitive context dependency. This is merely our *façon de parler*. What the analysis hopes to do is show how these properties *could* all have a common source. What their source truly is, that is a question which, strangely, has no answer.

Chapter 2

Context Change and Vague Predicates

The Basic Analysis

2.1 Chapter Overview

Adopting the program of study submitted in Section 1.5, our over-arching research goal is to understand the conglomeration of properties found peculiar to vague expressions as issuing from a single, common feature. With this goal in mind, the ultimate product of this thesis will be an analysis of the semantics of vague expressions which predicts the close relationship, discussed in Section 1.1, between an expression's having borderline cases, its susceptibility to sorites paradoxes, and its being context-dependent. Ideally, such an analysis would apply straightforwardly to all classes of vague expressions. In this thesis, however, we will nearly always confine our attention to vague *predicates.* I should point out, however, that our narrowed focus here is really quite unexceptional for this subject. The literature on vagueness predominantly discusses the semantics of vague predicates, since it is often correctly assumed by theorists that the standard machinery of semantic theory allows one to reduce the vagueness of all other semantic types to the ability for predicates to be vague [1].

The primary result of this thesis, then, is more accurately advertised as a semantic analysis of vague *predicates* which correctly predicts the relationship between these three targeted properties. In this chapter, we will begin moving towards that result by investigating the link between standard-sensitive context dependency and a predicate's having borderline cases. We will find that a possible source for this link is easily spotted once we have in place a semantic analysis for vague predicates using the tools of dynamic semantics. Dynamic

[1] For example, vague sentence radicals may be regarded as vague predicates of events, and vague determiners may inherit their vagueness from a vague predicate restricting their scope.

semantics will be our framework of choice throughout this thesis, and so a short refresher concerning its basic principles will be included in the next section.

After this refresher, we wade into our proposed analysis by first reviewing the "static semantics" usually assumed for gradable adjectives. This semantics captures the standard-sensitive context dependency of these vague expressions, and we thus propose that its basic form be extended to all vague predicates.

In Section 2.4, we will consider some simple data concerning the effect of assertions involving vague predicates on the information contained within a context. This data will inform our dynamic analysis of vague predicates, which is put forth in Section 2.5, and it is shown that this semantics correctly captures the data presented in 2.4.

Section 2.6 lays bare the link between standard-sensitive context dependency and the ability for a predicate to have borderline cases. The next section too-briefly discusses another interesting theoretical issue: the relationship between vagueness and adjectival comparison.

The last sections of this chapter spell out in greater detail the logic of vague predicates which our semantics yields. Comparisons are briefly made to the "supervaluational" semantics for vague predicates and logic it provides. Finally, in Section 2.9, the semantics of the vagueness operator "D" are introduced and a rather trivial consequence of adding this operator to our language is pointed out.

2.2 Dynamic Semantics: An Elementary Refresher

Dynamic semantics models the meaning of an expression as the effect which an utterance of that expression has on the information state of speakers engaged in a discourse. The philosophical fall-out of this is, of course, that dynamic semantics takes as the most basic semantic property of an expression its ability to alter the information available to a speaker. Whether or not one is comfortable with this metaphysics, the practical benefits of the framework are that it allows us a means of studying and describing in an explicit manner the effect particular assertions have on the context, and it provides us a way of modelling how utterances affect one another's meaning.

So much for the PR; now to the formalism. There are, in fact, a few distinct formalisms assembled under the rubric of "dynamic semantics". We will here first introduce what may be called the "Stalnakerian Approach"[2]; another formalism will be introduced in Chapter 3. An information state σ is formalized as a set of sequences of indices, in the simplest case as a set of possible worlds. Instead of "information state," we might sometimes also use the equivalent terms "knowledge state" or "context." Each information state is something like a time-slice of a full conversation, and the sequences making up the state are intended to indicate what the speakers of the conversation know or believe at that point of the discussion by each representing a "possibility" left open by what has been asserted thus far in the conversation. That is, speakers with

[2]See, for example, Stalnaker 1978.

the knowledge state σ will be said to "know" or to "believe" that ϕ if every sequence within σ models ϕ. A lower number of indices within a state, then, represents that less "possibilities" have been left open by the assertions making up the conversation, and hence more information has been made available to the speakers in the discourse. This prompts the definition that an information state σ' is more informative than an information state σ if $\sigma' \subset \sigma$.

Sentences are assigned as their semantic value a function from information states to information states, usually dubbed an "information change potential" or a "context change potential." This function represents the effect an utterance of that sentence has on the knowledge of the speakers in the discourse. In the simplest case, it may be characterized as merely the intersection of the input state with the proposition classically taken to be the semantic value of the sentence. For example, if our knowledge states consist of sets of sequences of possible worlds and variable assignments, we might model the meaning of the sentence "He$_i$ is a bad cook," $[\![\mathbf{bad}(\mathbf{cook}(x_i))]\!]$, as the function f defined as follows.

$$f : \wp(W \times G) \rightarrow \wp(W \times G)$$
$$f : \sigma \mapsto \sigma \cap \{\langle w, g \rangle : \langle w, g \rangle \models \mathbf{bad}(\mathbf{cook}(x_i))\}$$

The application of the information change potential assigned to a sentence ϕ to an input state σ is often called the "update" of σ with the proposition or information expressed by ϕ.

The analyses we propose in this thesis will not employ information change potentials of any complexity greater than that illustrated above. Therefore, within this thesis, the effect of updating any context σ with the information expressed by a sentence will either be a context σ' more informative than σ or will simply be σ. To prove that the utterance of ϕ within σ is informative, then, it will be sufficient to prove that $[\![\phi]\!](\sigma) \neq \sigma$.

To the uninitiated, it may not be obvious here what is to be gained by this approach. A neophyte may wonder why we do not simply take the meaning of the sentence in the standard way as the set of indices $\{\langle w, g \rangle : \langle w, g \rangle \models \mathbf{bad}(\mathbf{cook}(x_i))\}$ itself and then separately build a theory of information change around this. Indeed, it should later become clear to the reader that core content of our claims concerning vague language do not depend upon meanings being understood as "information change potentials". We use this framework because it offers a handy and elegant way of treating both the semantics and the pragmatics of utterances employing vague langauge, and this, the reader will find, *is* essential to the statement of our claims.

If these few, frantic comments do not sufficiently jog the reader's memory, or if the reader is new to the framework of dynamic semantics, it is suggested that he consult any of the following helpful sources: Dekker 1993; Groenendijk & Stokhof 1991; Groenendijk & Stokhof 1999; Groenendijk, Stokhof & Veltman 1996; Stalnaker 1978; van Benthem 1991; Veltman 1996.

2.3 A Classical Semantics for Vague Predicates

To find our way into a reasonable dynamic semantics for vague predicates, let us first recall what is standardly said about the semantics of gradeable adjectives used as predicates. The following is an abstraction and simplification of ideas found in the works of Cresswell 1976, Kennedy 1997, Seuren 1978, and von Stechow 1984, among others.

When used as predicates, gradable adjectives are undoubtedly vague. However, in many works on the semantics of gradable adjectives, the author begins by "bracketing" the issue of their vagueness. That is, the vagueness of these predicates is typically considered beyond the scope of such studies [3], and the simplifying assumption is made that gradable adjectives are not vague. I will thus follow suit, but the reader can be assured that their vagueness will re-enter the arena soon.

What can we say about the semantics of the sentence "John is tall", given this simplified picture? The analysis typically offered would seek to capture the standard-sensitive context dependency of this predicate by proposing the existence of a "tallness scale", a linearly ordered infinite set of "tallness degrees." These tallness degrees and their ordering relation are usually taken as primitive, though some analyses attempt to "construct" these entities from classes of objects or contexts [4]. It is then hypothesized that "John is tall" is true if and only if the degree on the tallness scale assigned to John is greater than the "standard" degree of tallness provided by the context [5].

Getting down to the formalism, we assign to "John is tall" the following truth conditions.

$$w, \delta_{tall} \models \text{tall(john) iff } f^{tall}(w)(john) >_{tall} \delta_{tall}$$

The analysis we adopt, then, is that "John is tall" is true at a world w and some adopted standard δ_{tall} iff John's tallness at w is greater than δ_{tall}. John's tallness at w is represented by the formula $f^{tall}(w)(john)$, where f^{tall} is a function from worlds and individuals to tallness degrees. Intuitively, this function takes an individual to its degree of tallness at w. Given this truth definition, it is obvious how we should treat the proposition that John is tall. $^{\cap}$tall(john) is simply the set of all pairs $\langle w, \delta_{tall}\rangle$ such that $w, \delta_{tall} \models$ tall(john).

Although analyses of this form are usually aimed at gradable adjectives, there is no reason why this approach could not be generalized to all predicates with standard-sensitive context dependency. That is, we propose that for any

[3] The studies I have in mind are primarily working towards a compositional semantics for the comparative.

[4] If the reader is skeptical of our taking such "degrees" as primitive, we will in the next chapter discuss some points in its favor.

[5] It should be pointed out that this semantics is neutral on the issue of whether the "standards" required for the interpretation of standard-sensitive predicates are "comparison classes" or something distinct from the implicit comparison class. The postulation of this contextually determined degree, for example, could be understood as abstracting from all the details going into the calculation of what might be expected from a member of the implicit comparison class.

standard-sensitive predicate P, there exists an associated infinite set of degrees O^P, linearly ordered without beginning or end point by a relation $<_P$, and an associated function f^P from pairs of worlds and objects to O^P [6] [7]. This scale and function are then used to interpret P in the expected way. Where $\delta \in O^P$ and α is some object in the domain of w, we offer the following truth conditions.

$$w, \delta \models P(\alpha) \text{ iff } f^P(w)(\alpha) >_P \delta$$

Similarly, the proposition ${}^{\cap}P(\alpha)$ is simply the set $\{\langle w, \delta\rangle : w, \delta \models P(\alpha)\}$, or equivalently, $\{\langle w, \delta\rangle : f^P(w)(\alpha) >_P \delta\}$.

The reader may worry that this definition assumes our language to include only one standard-sensitive context dependent predicate. How should we interpret the sentence $P(\alpha) \wedge Q(\beta)$, where both P and Q are standard-sensitive? One solution is obvious [8], but throughout this thesis we will make the simplifying assumption that our language has only one standard-sensitive predicate, usually written P. Due to the relationship between standard-sensitive context dependency and vagueness, this will mean that we likewise assume throughout that P is our language's only vague predicate.

The reader may also be reluctant about extending the notion of a linearly ordered scale of degrees to other standard-sensitive predicates besides gradable adjectives, such as nouns. Indeed, some theorists have expressed remorse over even extending this analysis to "multi-dimensional" gradable adjectives like "nice." If this reluctance arises purely from our straightforward appeal to the "degrees" of a predicate, the reader is advised to check his reluctance until having read our discussion in Section 3.5 supporting such an approach. However, if the reader accepts degrees as representations of the standards used in interpreting gradable adjectives, it is unclear why he should be in principle unwilling to accept degrees as representations of the standards used in interpreting nouns.

The source of the reader's reluctance here could also be the fact that we assume the degrees of a predicate to be in a *linear* ordering. Are the degrees of redness or of fatherliness to be found linearly ordered? To mitigate this worry, it may be said that the standards of these predicates certainly appear to be *ordered* by some measure of magnitude. One speaks informally of a person's "measure of niceness" or "level of fatherliness" being greater than another's. Our assumption that the ordering is linear merely allows a simpler formal model of how the degree to which a predicate holds of an object interacts with the contextually

---

[6] As in many philosophical works on the analysis of vague predicates, the formal metalanguage which we employ throughout this thesis will not be explicitly defined. The reader may understand this language to include nothing more than the logical constants and syntax available from first order logic. Moreover, at every sequence within a state, unless a sentence is of the form $P(\alpha)$ where P is a standard-sensitive predicate, a sentence is interpreted exactly as in a standard first order model structure.

[7] Regarding the lack of beginning or end point for $<_P$, this is intended to capture the intuition that if an object has a property to any degree, then it is possible for that object to have the property to a greater or a lesser degree. This intuition, however, requires that objects entirely lacking a property cannot have that property to any degree.

[8] Extend the sequence of indices to include degrees for all the vague predicates in the language.

given standard for the predicate to determine the truth of an utterance. The reader will find that none of our core claims regarding the nature of vagueness rest on the ordering of these standards being linear. Our claims *will* rest on the existence of such an ordering of standards, and on this ordering having certain other properties, but its linearity is for now only a matter of convenience. The reader may, if he likes, scoff at the very idea.

Finally, because of the correlation between standard-sensitive context dependency and vagueness, the classic semantics offered above for standard-sensitive predicates is thereby our proposed classic analysis for all vague predicates.

2.4 Vague Predicates and Information Change

As a final step in motivating our dynamic semantics for vague predicates, let us now consider some data concerning the manner in which assertions using vague predicates affect the information state of a speaker. We will here describe at a very rough and intuitive level the effect which an utterance of the sentence "John is tall" has on a speaker's knowledge state [9]. We focus our attention on the predicate "tall" because, as a "unidimensional" gradable adjective, the typical consequences of such utterances are especially vivid. Nevertheless, the reader should note that parallel facts do hold for all vague predicates [10].

Now, as with all context dependent expressions, the information which the assertion that John is tall provides the listener depends upon what the listener already knows or assumes about the context. With this in mind, four possible scenarios present themselves.

Case 1. Total Ignorance In this scenario, the listener knows or assumes absolutely nothing about what counts as "tall" in the context and knows or assumes absolutely nothing about John's height. When this is the case, the sentence "John is tall" seems to provide the listener with almost no information at all. The height of John has not in any way been narrowed down, and neither has the identity of the contextual parameter δ_{tall} .

To clarify, we are ignoring here the possible "default interpretation" of the sentence, "John is tall for a man". We imagine for the moment a case of complete ignorance, in which the listener does not make any judgement about what the speaker has in mind as his standard for tallness. Now, usually the listener does engage in a modicum of Gricean reasoning to arrive at some information about what is the contextual standard for "tall". Our interests here, however, demand that we abstract away from this and consider only the information provided by what the speaker's utterance "says." Without these usual, pragmatic

[9]Very similar data are discussed in Barker 2002. Moreover, the dynamic semantics for vague predicates proposed in the next section very closely resembles that put forth in Barker 2002. These facts and analyses were discovered independently, and the reader is strongly encouraged to consult Barker 2002 for a distinctly different discussion and argumentation.

[10]For example, if one were to say "Hand me that chair" while pointing to a stool, knowledge of the standard being employed for the predicate "chair" is affected much as in our Case 2.

assumptions, the listener cannot tell if the speaker means “tall for an ant,” “tall for a 5^{th} grader” or “tall for a building.” Moreover, since the standard cannot be determined from the content of the utterance, neither can any information about John’s height be obtained.

Case 2. Ignorance of Context, Awareness of Fact In this scenario, the listener knows or assumes nothing about what counts as “tall’ in the context but does know or assume something about John’s height. When the listener is in this state, the assertion that John is tall informs him only of the contextual standard for “tall.” Suppose that the listener knows an upper bound for John’s height; John must be less than seven feet tall. The listener can then deduce that the contextual standard for “tall” must be below seven feet. However, given no prior knowledge of the tallness standard, the listener cannot conclude anything regarding John’s height.

Case 3. Total Awareness In this scenario, the listener knows or assumes something both about the standard for “tall” and about John’s height. It now seems that the listener could be in a position to either deny or agree to an assertion of John’s being tall. If the listener already knows that John is above seven feet and that the contextual standard for tall is below seven feet, then he will be ready to nod his head to any assertion that John is tall. Moreover, this assertion could also inform the listener simultaneously of John’s height and the contextual standard for tallness. Suppose that we know Sue is taller than John, that Sue is 5 feet, and that Mary is short. If we learn that John is tall, we simultaneously learn that the value for tallness is below 5 feet and that John is taller than Mary.

Case 4. Awareness of Context, Ignorance of Fact In this scenario, the listener assumes or knows something about what counts as tall, but he doesn’t assume or know anything about John’s tallness. The assertion that John is tall provides the listener with information regarding John’s height, but nothing is learned about the contextual standard for “tall.” If the listener knows that the standard for “tall” must be above six feet, then he has learned that John is at least six feet tall.

2.5 A Dynamic Semantics for Vague Predicates

How might we capture these processes of information change within a dynamic semantics for vague predicates? Let P be a vague predicate and O^P be its associated set of degrees. Keeping our static semantics in mind, suppose that we take a state of total ignorance on the part of the listener to be the full set $W \times O^P$. Now, given such a loathsome dearth of knowledge, we add to the needy listener’s information by asserting sentences ϕ. The effect of this will be the intersection of our information state with the set of world-degree pairs constituting the proposition expressed by ϕ. Clearly, then, the dynamic semantics we will propose is simply one in which the information state of a discourse employing the vague predicate P is taken to be a subset of $W \times O^P$, and in which the propositions expressed by utterances of the form $P(\alpha)$ are as was proposed in Section 2.3. For example, the meaning of the sentence "John is tall" is the function $[\![\text{tall(john)}]\!]$ defined as follows.

$$[\![\text{tall(john)}]\!] : \wp(W \times O^{tall}) \to \wp(W \times O^{tall})$$
$$[\![\text{tall(john)}]\!] : \sigma \mapsto \sigma \cap \{\langle w, \delta_{tall}\rangle : \langle w, \delta_{tall}\rangle \models \text{tall(john)}\}$$

Let us see how this proposal coheres with our data from the previous section. To neaten the appearance of our arguments, however, let us first introduce the following definitions.

Definition 5.1 Let $\sigma \subseteq W \times O^P$ be an information state. The **degrees available at** σ is the set $\{\delta : \exists w \in W \ \langle w, \delta\rangle \in \sigma\}$. The **worlds available at** σ is the set $\{w : \exists \delta \in O^P \ \langle w, \delta\rangle \in \sigma\}$.

1. Total Ignorance Suppose that we add the proposition $^{\cap}$tall(john) to our state of total ignorance $\sigma = W \times O^{tall}$. The resulting state σ' will be a proper subset of σ, and so in that sense the assertion was informative. However, σ' is still such that for every $w \in W$ there is a pair $\langle w, \delta_{tall}\rangle \in \sigma'$ and for every $\delta_{tall} \in O^{tall}$ there is a pair $\langle w, \delta_{tall}\rangle \in \sigma'$.

Proposition 5.2 For every $w \in W$ there is a pair $\langle w, \delta_{tall}\rangle \in \sigma'$.

PROOF: Suppose $w \in W$ and $f^{tall}(w)(john) = \delta'_{tall}$. Adding $^{\cap}$tall(john) to a state σ removes all pairs $\langle w, \delta_{tall}\rangle$ such that $\delta'_{tall} \leq_{tall} \delta_{tall}$. But, if σ is $W \times O^{tall}$, then since $\langle O^{tall}, <_{tall}\rangle$ has no beginning point, there is a pair $\langle w, \delta_{tall}\rangle \in \sigma$ such that $\delta'_{tall} >_{tall} \delta_{tall}$. So $\langle w, \delta_{tall}\rangle$ survives the update and is an element of the resulting state σ'.

Proposition 5.3 For every $\delta_{tall} \in O^{tall}$ there is a pair $\langle w, \delta_{tall}\rangle \in \sigma'$.

PROOF: Similar to that for the previous proposition.

Therefore, the listener still does not have any information about the identity of the height of John or of the contextual standard for "tall." No possible worlds w or possible standards δ_{tall} have been ruled out by the assertion. All the listener learns from the utterance is that whatever John's height is, it will be greater than whatever the standard for "tall" turns out to be. Moreover, this information is reflected in the resulting state σ', since for every pair $\langle w, \delta_{tall}\rangle \in \sigma'$, $f^{tall}(w)(john) >_{tall} \delta_{tall}$.

2. Ignorance of Context, Awareness of Fact Let us now suppose that our initial knowledge state σ is one in which an upper bound on John's height is known, but the listener lacks *completely* any information about the contextual standard for "tall." That is, this listener does not have any information as to the identity of this contextual parameter, nor any information linking its identity to factual states of affairs. σ, then, is the full product $W' \times O^{tall}$, where $W' \subset W$ and $\exists \delta_{tall} \in O^{tall}\ \forall w \in W'\ f^{tall}(w)(john) \leq_{tall} \delta_{tall}$. If we add the information that John is tall, the resulting state σ' includes no new information about the world.

Proposition 5.4 The set of worlds available at σ is the same as the set of worlds available at σ'.

PROOF: Suppose that w is in the set of worlds available at σ and $f^{tall}(w)(john) = \delta'_{tall}$. Since $\sigma = W' \times O^{tall}$, and $\langle O^{tall}, <_{tall}\rangle$ is without beginning point, there is a pair $\langle w, \delta_{tall}\rangle \in \sigma$ such that $\delta'_{tall} >_{tall} \delta_{tall}$. Therefore, $\langle w, \delta_{tall}\rangle \in \sigma'$ and so w is a world available at σ'. Due to the nature of updates, the worlds available at σ must then be equal to the worlds available at σ'.

However, since an upper bound is known for John's height, the assertion does supply information about the standard for "tall."

Proposition 5.5 The set of degrees available at σ does not equal the set of degrees available at σ'.

PROOF: $\exists \delta_{tall} \in O^{tall}\ \forall w \in W'\ f^{tall}(w)(john) \leq_{tall} \delta_{tall}$. Suppose that $\forall w \in W'\ f^{tall}(w)(john) \leq_{tall} \delta'_{tall}$. Consider an arbitrary pair $\langle w, \delta'_{tall}\rangle$ where $w \in W'$. Since $f^{tall}(w)(john) \leq_{tall} \delta'_{tall}$, $f^{tall}(w)(john) \not>_{tall} \delta'_{tall}$, and so $w, \delta'_{tall} \not\models$ tall(john). Therefore, $\langle w, \delta'_{tall}\rangle$ would not survive the update with $^{\cap}$tall(john), and so $\langle w, \delta'_{tall}\rangle \notin \sigma'$. Since w was arbitrary, δ'_{tall} is not in the set of degrees available at σ'. By definition, however, $\delta'_{tall} \in O^{tall}$; thus, δ'_{tall} is in the set of degrees available at σ.

The result is that the upper bound for John's height now places an upper bound on the contextually given standard for tall.

Proposition 5.6 There exists a $\delta'_{tall} \in O^{tall}$ such that for all $\langle w, \delta_{tall}\rangle \in \sigma'$, $\delta_{tall} <_{tall} \delta'_{tall}$.

PROOF: $\exists \delta_{tall} \in O^{tall}\ \forall w \in W'\ f^{tall}(w)(john) \leq_{tall} \delta_{tall}$. Suppose that $\forall w \in W'\ f^{tall}(w)(john) \leq_{tall} \delta'_{tall}$. Let $\langle w, \delta_{tall}\rangle \in \sigma'$. $w, \delta_{tall} \models$ tall(john). Therefore, $f^{tall}(w)(john) >_{tall} \delta_{tall}$ and so $\delta_{tall} <_{tall} \delta'_{tall}$. Since $\langle w, \delta_{tall}\rangle$ was arbitrary, for any $\langle w, \delta_{tall}\rangle \in \sigma'$, $\delta_{tall} <_{tall} \delta'_{tall}$.

Thus, the listener receives no information about the world, but he does obtain information regarding the standard employed for "tall".

3. Total Awareness In our third scenario, the listener has information regarding both John's height and the standard used for "tall". This case is actually the most often occurring, it being a rather fanciful state of affairs where the listener bears no assumptions regarding the contextual standard. Typically, when a speaker asserts that John is tall, the listener accommodates the utterance so that their information state includes the information that the standard for "tall" is something plausible for human beings. The consequence is that the information that John is tall is never added to a state of complete ignorance regarding what counts as tall, and so some information about the world is always attained.

Now, it was said that in this third scenario, it should be possible for a speaker to already believe that John is tall. We would with our adopted formalism represent a state of total awareness σ as some subset of $W' \times O_i^{tall}$, where $W' \subset W$ and $O_i^{tall} \subset O^{tall}$. Clearly, one such σ could be a state in which for every $\langle w, \delta_{tall}\rangle \in \sigma$, $f^{tall}(w)(john) >_{tall} \delta_{tall}$. Since every pair in this putative state models "tall(john)", we thus correctly capture that there are some states of total awareness in which the speakers already believe that John is tall. Similarly, we predict that other states of total awareness may be ones in which speakers already believe that John is not tall.

Furthermore, we observed in Section 2.4 that a listener in a state of total awareness may be in a position where the information that John is tall is helpful, and informs him both of John's height and of the standard for "tall." Suppose that our listener knows that Sue is taller than John, that Sue is 5 feet tall and that Mary is not tall. If we add to this state the proposition $^\cap$tall(john), the resulting state σ' contains both less worlds and less tallness degrees.

Proposition 5.7 Let $\sigma \subseteq W' \times O_i^{tall}$ where $W' \subset W$ and $O_i^{tall} \subset O^{tall}$. Let σ contain only the information that Sue is taller than John, that Sue is 5 feet tall and that Mary is not tall. The update of σ with $^\cap$tall(john) results in a state σ' such that the set of worlds available at σ does not equal the set of worlds available at σ'.

PROOF: Consider the world w in the set of worlds available at σ in which John is shorter than Mary; i.e., $f^{tall}(w)(john) <_{tall} f^{tall}(w)(mary)$. Let $\langle w, \delta_{tall}\rangle$ be a pair in σ. Since $\langle w, \delta_{tall}\rangle \not\models$ tall(mary), $f^{tall}(w)(mary) \leq_{tall} \delta_{tall}$, and so $f^{tall}(w)(john) <_{tall} \delta_{tall}$, and thus $\langle w, \delta_{tall}\rangle \not\models$ tall(john). Therefore, $\langle w, \delta_{tall}\rangle$ does not survive the update with $^\cap$tall(john) and so $\langle w, \delta_{tall}\rangle \notin \sigma'$. Since $\langle w, \delta_{tall}\rangle$ was arbitrary, w does not survive in σ', and so w is not in the set of worlds available at σ'.

Proposition 5.8 Let σ be as is stated in Proposition 5.7. The update of σ with $^\cap$tall(john) results in a state σ' such that the set of degrees available at σ

does not equal the set of degrees available at σ'.

PROOF: Consider the tallness degree δ_{tall} in the set of degrees available at σ such that $\delta_{tall} = 5ft$ [11]. Let $\langle w, \delta_{tall}\rangle$ be an element of σ. Since $f^{tall}(w)(sue) >_{tall} f^{tall}(w)(john)$ and $f^{tall}(w)(sue) = \delta_{tall}$, $f^{tall}(w)(john) <_{tall} \delta_{tall}$. Therefore, $\langle w, \delta_{tall}\rangle \not\models$ tall(john), and so $\langle w, \delta_{tall}\rangle \notin \sigma'$. Since $\langle w, \delta_{tall}\rangle$ was arbitrary, δ_{tall} does not survive in σ', and so δ_{tall} is not in the set of degrees available at σ'.

The reader is also invited to confirm that we may with this dynamic semantics predict that in the resulting state σ' the speakers know that John is taller than Mary, that the standard for tallness is below five feet, that Sue is tall, and that Mary is below five feet in height.

4. Awareness of Context, Ignorance of Fact Finally, let us suppose that our initial knowledge state σ is one in which a lower bound on the contextual standard for "tall" is known, but the listener lacks *completely* any information regarding the vertical dimension of John. This listener, then, does not have any information regarding the height of John, nor any information linking his height to the identity of the standard for "tall." σ, is therefore the full product $W' \times O_i^{tall}$ where $W' \subseteq W$, $O_i^{tall} \subset O^{tall}$, $\exists \delta'_{tall} \in O^{tall}\ \forall \delta_{tall} \in O_i^{tall}\ \delta_{tall} >_{tall} \delta'_{tall}$ and $\forall \delta_{tall} \in O^{tall}, \exists w \in W'$ such that $f^{tall}(w)(john) = \delta_{tall}$. If we add the information that John is tall, the resulting state σ' includes no new information about the standard for "tall."

Proposition 5.9 The set of degrees available at σ is the same as the set of degrees available at σ'.

PROOF: Suppose that δ_{tall} is in the set of degrees available at σ. Since $\sigma = W' \times O_i^{tall}$, there is a pair $\langle w, \delta_{tall}\rangle \in \sigma$ such that $f^{tall}(w)(john) >_{tall} \delta_{tall}$. Therefore, $\langle w, \delta_{tall}\rangle$ survives update with $^{\cap}$tall(john) and so $\langle w, \delta_{tall}\rangle \in \sigma'$. Thus, δ_{tall} is in the set of degrees available at σ', and by the nature of update, the degrees available at σ and at σ' must be the same.

However, since there exists a lower bound for the tallness standard, such an assertion does supply information about John's height.

Proposition 5.10 The set of worlds available at σ does not equal the set of worlds available at σ'.

PROOF: $\exists \delta'_{tall} \in O^{tall}\ \forall \delta_{tall} \in O_i^{tall}\ \delta_{tall} >_{tall} \delta'_{tall}$. Therefore, $\exists w \in W'\ \forall \delta_{tall} \in O_i^{tall}\ f^{tall}(w)(john) <_{tall} \delta_{tall}$. Suppose that $\forall \delta_{tall} \in O_i^{tall}\ f^{tall}(w')(john) <_{tall} \delta_{tall}$. Consider an arbitrary pair $\langle w', \delta_{tall}\rangle$, where δ_{tall} is a degree available at σ. Since $f^{tall}(w')(john) <_{tall} \delta_{tall}$, $w', \delta_{tall} \not\models$ tall(john). Therefore, $\langle w', \delta_{tall}\rangle$ would not survive the update with $^{\cap}$tall(john), and so $\langle w', \delta_{tall}\rangle \notin \sigma'$. Since

[11] That is, this tallness degree is assigned to any object five feet in height.

$\langle w', \delta_{tall}\rangle$ was arbitrary, w' is not a world available at σ'. By definition, however, w' is a world available at σ.

The result is that the lower bound for the standard of tallness now places a lower bound on John's height.

Proposition 5.11 There exists a $\delta'_{tall} \in O^{tall}$ such that for all $\langle w, \delta_{tall}\rangle \in \sigma'$, $f^{tall}(w)(john) >_{tall} \delta'_{tall}$.

PROOF: Let δ'_{tall} be such that $\forall \delta_{tall} \in O_i^{tall}\ \delta_{tall} >_{tall} \delta'_{tall}$. Let $\langle w, \delta_{tall}\rangle \in \sigma'$. By definition, $w, \delta_{tall} \models$ tall(john). Therefore, $f^{tall}(w)(john) >_{tall} \delta_{tall}$ and so $f^{tall}(w)(john) >_{tall} \delta'_{tall}$. Since $\langle w, \delta_{tall}\rangle$ was arbitrary, for any $\langle w, \delta_{tall}\rangle \in \sigma'$ $f^{tall}(w)(john) >_{tall} \delta'_{tall}$.

Therefore, just as we informally observed in Section 2.4, although the listener receives no information about the standard for "tall", he does obtain information regarding John's height.

The proofs above demonstrate that the proposed dynamic semantics for vague predicates is at least a plausible way of representing the information contained in assertions employing those predicates and how such assertions affect the knowledge of a listener. Equipped with this semantics, we could, if we wished, pursue further this study of the conversational dynamics of vague language. Indeed, Kyburg & Morreau 2000 and Barker 2002 independently develop similar systems in their more focused investigations of the subtle interplay between speakers using vague predicates. Although we certainly recommend that our semantic analysis be used for such research, in this thesis our aim is a bit to the right of that target. Our curiosity is focused on how this analysis might pressed into providing us some insight into the relationship between standard-sensitive context dependency, having borderline cases, and being susceptible to sorites paradoxes. In the following section, we will begin indulging this curiosity by connecting our proffered dynamic semantics to the fact that all standard-sensitive context dependent predicates have the potential for borderline cases.

2.6 Truth and Borderlines

A distinctive feature of dynamic semantics is that it does not centrally concern itself with the notion of "truth." Instead, the effects of an utterance on the knowledge state of a speaker take pride of place. Since this concept of "information change" is taken as primitive, other notions, such as the "content" of an utterance, must be derived from the basic ability an utterance has to affect the context. Consequently, any statement within dynamic semantics of what "truth" for a particular language is will require this property to be in some way emergent from the state of the discourse and the information change potential of the utterance. However, there is already in dynamic semantics a concept which comes very close to, and may be profitably identified with truth: the notion of "support."

Definition 6.1 An information state σ **supports** a sentence ϕ if the update of σ with ϕ is σ; if $[\![\phi]\!](\sigma) = \sigma$. An information state σ **rejects** a sentence ϕ if the update of σ with ϕ is the empty set; if $[\![\phi]\!](\sigma) = \emptyset$.

Note that the definitions above entail that a state σ supports ϕ if and only if it rejects $\neg\phi$ and that a state σ supports $\neg\phi$ if and only if it rejects ϕ. Moreover, the notion of "support" comes quite close to the analyses of "truth" in pragmatist philosophies of science and in constructivist philosophies of mathematics. This nice harmony suggests that we identify the truth value of a formula ϕ at a state σ with the mode in which σ supports ϕ.

Definition 6.2 A sentence ϕ is **T** (true) at an information state σ if σ supports ϕ. A sentence ϕ is **F** (false) at an information state σ if σ rejects ϕ. A sentence ϕ is **I** (indeterminate) at an information state σ if ϕ is neither T nor F at at σ.

Our dynamic semantics for vague predicates combined with this definition of "truth" will thus provide us a three-valued logic of vague predicates. This three valued logic will be explored in greater depth in Section 2.8.

Our development of a multi-valued logic for vague predicates should, of course, not come as much of a surprise. Logics for vague predicates are typically multi-valued due to the existence of borderline cases. It is, in fact, often the case in such logics that an atomic sentence $P(\alpha)$ takes on a non-standard truth value precisely when α is a borderline case for P. In our logic, we will not have this close a correlation between indeterminacy of truth value and borderline cases, but there will be instances of sentences $P(\alpha)$ assigned the value I in which α may be profitably identified as a borderline case.

Recall that a "borderline case" α for a vague predicate P is an entity such that it is unknown whether $P(\alpha)$, and it seems that *no increase in the knowledge one has of the world* could determine whether $P(\alpha)$. Exactly these cases occur when the knowledge state σ contains, for example, the information that the standard for "tall" is between 6 and 5 feet and it is known that John's height is between 5' 8" and 5' 10". Intuitively, no increase in the knowledge of John's height can distinguish him as tall or not tall. One's inability to assert in such a context "John is tall" has not to do with an ignorance of fact, but with the fact that the context does not provide an absolute, precise boundary for "tall," and John's height is known to fall within the "gap" of heights permitted by the context. Finally, the indeterminacy of whether John is tall may only be resolved if the speakers within the discourse adopt a more exacting, precise concept of what "tall" means. That is, if the discussants place the possible standards for "tall" between 5' 8.5" and 5' 9", then John is no longer a borderline case of tallness. Although it may remain unknown whether John is tall, the speakers are in a position now to determine whether this is the case by increasing their knowledge of John's height.

These considerations prompt the following definition.

Definition 6.3 Let P be a vague predicate and O^P be its associated set of degrees. An object α is a **borderline case** for P at an information state $\sigma \subseteq W \times O^P$ if $P(\alpha)$ is I at σ, and $P(\alpha)$ is I at any information state σ' such that the worlds available at σ' are a subset of the worlds available at σ, the degrees available at σ' and σ are the same, and for any w in the set of worlds available at σ', if $\langle w, \delta \rangle \in \sigma$ then $\langle w, \delta \rangle \in \sigma'$.

Loosely speaking, we define a borderline case for P at an information state σ as an object which is indeterminately P and remains indeterminately P when we increase our information *only* regarding what worlds are possible, not what standards for P are possible nor how those standards relate to the available worlds.

As has already been pointed out, a "borderline case" by the definition in 6.3 bears a strong similarity to borderline cases as informally introduced in 1.1. Speakers within a context where α is a borderline case for P will not know whether $P(\alpha)$ or $\neg P(\alpha)$, and so by the basic tenets of dynamic semantics, they will not be able to assert whether $P(\alpha)$ or $\neg P(\alpha)$. Moreover, by our condition that borderline cases α for P at σ must remain indeterminately P at any state σ' constituting only an increase in world knowledge from σ, we capture the intuition that the uncertainty surrounding borderline cases for a predicate cannot be resolved by increasing one's knowledge of the world. That is, we capture the intuition that an object's being a borderline case for a predicate is due to an indeterminacy, an openness in the meaning of the predicate.

Certainly, then, it appears then that our formalization in 6.3 of the notion of a "borderline case" captures important properties of borderline cases as informally understood. The connection we claim to have discovered here between a predicate's having borderline cases and its being standard-sensitive is that standard-sensitive context dependency implies having borderline cases; for any standard-sensitive predicate P there exists a context σ and an object α such that α is a borderline case for P at σ.

Theorem 6.4 (Law of Borderlines) Let P be a standard-sensitive context dependent predicate, O^P be its associated set of degrees, and f^P be its associated function from worlds and objects to O^P. There exists a context $\sigma \subseteq W \times O^P$ and an object α such that α is a borderline case for P at σ.

PROOF:[12] Let α be some object, let $O^P_i \subset O^P$ be such that $O^P_i \neq \emptyset$ and for all $\delta \in O^P_i$, there is a $\delta' \in O^P_i$ such that $\delta' <_P \delta$, let $W' \subset W$ be such that

[12]The reader will observe that our argument actually establishes the stronger proposition that for any object α and standard-sensitive predicate P, there is an information state at which α is a borderline case for P. The claim that any object is a potential borderline for any vague predicate is, of course, not credible. This glitch, however, is merely the result of

$O_i^P = \{\delta : \exists w \in W', f^P(w)(\alpha) = \delta\}$ and $\sigma = W' \times O_i^P$. Let $\delta, \delta' \in O_i^P$ be such that $\delta' <_P \delta$. Now, let $w \in W'$ be such that $f^P(w)(\alpha) = \delta$. Clearly, $\langle w, \delta\rangle, \langle w, \delta'\rangle \in \sigma$. Moreover, since $f^P(w)(\alpha) >_P \delta'$ and $f^P(w)(\alpha) \not>_P \delta$, $P(\alpha)$ is I at σ.

Now let σ' be such that the worlds available at σ' are a subset of the worlds available at σ, the degrees available at σ' and σ are the same, and for any w in the set of worlds available at σ', if $\langle w, \delta\rangle \in \sigma$ then $\langle w, \delta\rangle \in \sigma'$. Let w be in the worlds available at σ', and $f^P(w)(\alpha) = \delta$. Since $w \in W'$ and $\sigma = W' \times O_i^P$, $\langle w, \delta\rangle \in \sigma$ and so $\langle w, \delta\rangle \in \sigma'$. Moreover, since there is a $\delta' \in O_i^P$ such that $\delta' <_P \delta$, we know that $\langle w, \delta'\rangle \in \sigma'$. Thus, since $f^P(w)(\alpha) \not>_P \delta$ but $f^P(w)(\alpha) >_P \delta'$, $P(\alpha)$ is I at σ.

Thus, α is a borderline case for P at σ.

We claim that our having demonstrated the Law of Borderlines establishes that our semantic theory predicts all standard-sensitive context dependent predicates to have borderline cases.

Now, although the definition in 6.3 is clear enough to offer a correct proof of Theorem 6.4, it is perhaps unclear to the reader what the underlying philosophy, or "theory," of borderline cases adopted in this thesis truly is. Indeed, the definition in 6.3 leaves much room to play on this issue. To a certain extent, we would like to remain officially agnostic on this question, leaving it to the reader to decide how best to describe our claims informally. One way of characterizing our theory of borderline cases in 6.3 is that α is a borderline case of a vague predicate P within a discourse σ when the speakers' use of P in σ underdetermines its meaning in such a way that $P(\alpha)$ is neither true nor false in that context, and *only* a change in the way the speakers use P can create a context in which $P(\alpha)$ would be true or false. Such a characterization would link our claims to that strand in the philosophical literature which describes borderline cases as the result of the way in which the data presented to a language learner must vastly underdetermine the meanings of the predicates in the language [13]. Perhaps another equally valid way of informally describing our account is that α is a borderline case of a vague predicate P within a discourse σ when the speakers within σ cannot determine which of a number of precise meanings P has. That is, α's being a borderline case of P is due to a kind of ignorance of P's precise meaning. Such a characterization would then link our claims to those "maverick" positions commonly dubbed "epistemic" theories of vagueness.

Within our framework, however, there is no intelligible difference between speakers within a context "failing to adopt a precise meaning for P" and their "being ignorant of the precise meaning of P". One is free, then, to adopt either of

---

our formal theory abstracting away from the manner in which the standards in a context are arrived at by computing over an implicit comparison class and, moreover, from the manner in which the set of possible comparison classes is constrained by what people are willing to admit as a "kind". On this latter point, see Graff 2000, Section III.

[13] It would also put our claims in harmony with the view that borderline cases are the result of an "openness" in the "semantic rules" governing the use of the predicate, if one understands those rules as appealing to the contextually varying standard for the predicate.

the characterizations above. Arguably, though, whatever is one's favorite theory of the nature of borderline cases, our formal definition in 6.3 may be projected onto it, and so one's philosophical persuasions do not undermine agreement that our semantic framework predicts all standard-sensitive context dependent predicates to admit of borderline cases.

2.7 Vagueness and Comparison

In this section, we will take a scenic detour from the main narrative of the chapter and swiftly cruise through some issues relating to the semantics of the comparative construction. Section 2.7.1 discusses the relationship between vagueness and gradability in adjectives, while Section 2.7.2 briefly addresses the interesting difference between the contextual effects of assertions employing the comparative construction and those employing the absolute form of a gradable adjective.

2.7.1 The Relationship Between Gradability and Vagueness

In Section 2.3, we explored some objections against extending the degree-based static semantics for gradable adjectives to all standard-sensitive context dependent predicates. However, one possible objection to our approach which was not introduced in that section is that it seems to draw too close a connection between the gradability and the vagueness of an adjective.

An adjective is said to be "gradable" if it is possible for that adjective to appear in the comparative form. For example, the adjectives "tall," "fat," and "interesting" are gradable, while "pregnant" and "socialist" are not. We noted in 2.3 that all gradable adjectives are standard-sensitive, and so therefore vague. However, it is *not* the case that all vague adjectives are gradable. "Dead," for example, is vague, in as much as it has imaginable borderline cases (zombies, vampires, cryogenically frozen individuals), is susceptible to a sorites paradox (imagine a sequence of terminal individuals, each with successively lower vital signs), and is to a certain degree standard-sensitive (as revealed by the distinction between "brain death" and "heart death"). The impossibility of the expressions "deader" and "more dead," though, reveals that the adjective is not gradable. This distinction between gradability and vagueness is further highlighted by the fact that nouns and verbs can never appear in a comparative form, even when they are monstrously vague. As Kamp 1975 points out, sentences such as "A barn is more (of) a building than an igloo is" or "Dave more smootched Ariel than he did Suzanne" are always somewhat awkward, and are never uncontroversially true except in cases in which one of the two comparands is clearly not within the extension of the predicate.

A certain problem, then, arises for our attempt to extend the degree-based semantic analysis of gradable adjectives to all other vague predicates. The use of degrees in stating the semantics of gradable adjectives is typically intended as a means of relating the semantics of these adjectives when used as predicates to the semantics of their comparative form. "Dave is taller than Jacob," for example, is analyzed as meaning "the degree of Dave's tallness is greater than the degree of Jacob's tallness." But, if all vague predicates have an associated degree ordering, and if all the comparative seems to require for its semantics is the existence of an associated degree ordering, then why cannot all vague predicates appear in the comparative?

Notice that one cannot as a solution to this problem take the line that there is simply a syntactic restriction of the comparative morpheme "er/more" to the class of adjectives. "Dead," after all, is a vague adjective which still does not naturally appear with the comparative morpheme. One might, however, still propose that the difference between gradable and non-gradable vague adjectives is to be found within the intricacies of the syntactic structure of the comparative. Kennedy 1997, for example, demonstrates that the most likely syntactic analyses of sentences containing the comparative and non-comparative forms of an adjective imply that one does best, in designing a compositional semantics for the comparative construction, to take gradable adjectives as having as their semantic value a function from worlds and objects to degrees. That is, if the arguments of syntacticians are to be believed, then the most elegant compositional semantics for the comparative requires the semantics of the lexical item "tall" to be, roughly speaking, the function f^{tall}.

Perhaps, then, the difference between gradable adjectives and all other vague predicates is simply one of logical type. Following the analysis of Kennedy 1997, the comparative morpheme can only take as its argument a function of the type $\langle\langle s, e\rangle, \delta\rangle$. However, non-gradable vague predicates, such as "dead" or "building," are simply functions of the usual predicational type $\langle e, t\rangle$. Thus, the inability for "dead," "socialist," and "in New Jersey" to appear in the comparative, despite their having an associated ordering of degrees, is due to a straightforward type mismatch.

At the moment, this may sound like nothing more than an unenlightening, *ad hoc* stipulation. Kennedy's work, however, may reveal it to have some unexpected empirical bite. For example, Kennedy 1997 independently argues that degree modifiers such as "very" take as their arguments functions from worlds and objects to degrees. Our proposed type distinction, then, may explain the impossibility of the expressions "very dead" and "very in New Jersey." For the moment, though, this remains largely an issue for further research.

2.7.2 The Effects of the Comparative upon the Context

Barker 2002 (p. 13) claims that, unlike sentences in which a gradable adjective is used as a predicate, the utterance of a sentence containing only the comparative form a gradable adjective cannot affect the listener's knowledge of the contextual standard for the adjective. That is, in no situation does an utterance of "Dave is taller than Phil" inform one of the standards being adopted for "tall." The interpretation of the sentence does not rely on those standards, and so no utterance of the sentence affects one's knowledge of them.

Barker 2002 attempts to capture this feature of the comparative by showing that when two common analyses of the comparative are transported into his dynamic semantics for gradable adjectives, both predict that assertions containing only the comparative form of the adjective should not affect what degrees are available within the context. These two common analyses, however, are rather forcibly criticized in Kennedy 1997. Furthermore, contrary to Barker's claims, an assertion of the sentence "Dave is taller than Phil" *can* sometimes affect the listener's knowledge of the contextual standard for "tall." Imagine that I have told you that Dave is six feet tall and that Phil is tall. So far, you have no knowledge of the standard for "tall." But, if I tell you that Dave is taller than Phil, you thereby become informed that the standard for "tall" is below six feet. Read charitably, however, Barker 2002 assumes that the state in which the assertion is made is one in which the listener lacks *completely* any information regarding the identity of the standard for "tall" or how that standard relates to factual states of affairs. In such cases, it does seem that assertions of the sentence "Dave is taller than Phil" cannot affect a listener's knowledge of the standard being used for "tall."

Let us see, then, whether by transporting the analysis of Kennedy 1997 into our framework, we cannot more accurately predict the differing information change potential of the comparative. Kennedy 1997 proposes that sentences such as "Dave is taller than Phil" be given the following truth conditions.

$$w, \delta_{tall} \models \text{more(tall)(dave,phil) iff } f^{tall}(w)(dave) >_{tall} f^{tall}(w)(phil)$$

Clearly, then, if P is a standard-sensitive predicate and O^P is its associated set of degrees, our dynamic semantics for sentences of the form $more(P)(\alpha, \beta)$ will be the following.

$$[\![more(P)(\alpha, \beta)]\!] : \wp(W \times O^P) \to \wp(W \times O^P)$$
$$[\![more(P)(\alpha, \beta)]\!] : \sigma \mapsto \sigma \cap \{\langle w, \delta \rangle : f^P(w)(\alpha) >_P f^P(w)(\beta)\}$$

With this dynamic semantics in place, we may show that some assertions of "Dave is taller than Phil" can supply information regarding the contextual standard for tall.

Proposition 7.1 Let σ be a state containing only the information that Dave is six feet tall and Phil is tall, and let σ' be the state $[\![\text{more(tall)(dave,phil)}]\!](\sigma)$. The set of degrees available at σ does not equal the set of degrees available at σ'.

PROOF: Let δ_{tall} be the degree available in σ such that $\delta_{tall} = 6ft$[14]. Let $\langle w, \delta_{tall} \rangle$ be a pair in σ. $\langle w, \delta_{tall} \rangle$ must model tall(phil), and so $f^{tall}(w)(phil) >_{tall} \delta_{tall}$. Similarly, $\langle w, \delta_{tall} \rangle$ must model six-feet(dave), and so $f^{tall}(w)(dave) = \delta_{tall}$. Thus, $f^{tall}(w)(phil) >_{tall} f^{tall}(w)(dave)$, and so $\langle w, \delta_{tall} \rangle$ does not survive the update with $^\cap$more(tall)(dave,phil), and so $\langle w, \delta_{tall} \rangle \notin \sigma'$. Since $\langle w, \delta_{tall} \rangle$

[14]See footnote 11 for an explanation of this notation.

was arbitrary, δ_{tall} is not a degree available at σ'.

Moreover, we may show that such informative assertions of "Dave is taller than Phil" can supply an upper bound for the contextual standard.

Proposition 7.2 Let σ' be as stated in Proposition 7.1. If δ_{tall} is a degree available at σ' and $\delta'_{tall} = 6ft$, then $\delta_{tall} <_{tall} \delta'_{tall}$.

PROOF: Suppose, for a contradiction, that $\langle w, \delta_{tall}\rangle \in \sigma'$ but $\delta_{tall} \geq_{tall} \delta'_{tall} = 6ft$. Since $\langle w, \delta_{tall}\rangle$ must model tall(phil), $f^{tall}(w)(phil) >_{tall} \delta_{tall}$, and so $f^{tall}(w)(phil) >_{tall} \delta'_{tall}$. However, $\langle w, \delta_{tall}\rangle$ must also model six-feet(dave), and so $f^{tall}(w)(dave) = \delta'_{tall}$. Thus, $f^{tall}(w)(phil) >_{tall} f^{tall}(w)(dave)$, which contradicts the assumption that $\langle w, \delta\rangle \in \sigma'$. Therefore, $\delta_{tall} <_{tall} \delta'_{tall}$.

Finally, we may with this analysis prove that in contexts containing no information regarding the identity of the contextual standard or how that standard relates to possible states of affairs, felicitous assertions of "Dave is taller than Phil" cannot provide any information regarding the contextual standard for "tall."

Proposition 7.3 Let σ be equal to $W' \times O^{tall}$, where $W' \subseteq W$, and let $\sigma' \neq \emptyset$ be the state $[\![\text{more(tall)(dave,phil)}]\!](\sigma)$. The set of degrees available at σ and the set of degrees available at σ' are the same.

PROOF: Suppose that δ_{tall} is a degree available at σ, and that $w \in W'$ is a world available at σ'. By definition, $\langle w, \delta_{tall}\rangle \in \sigma$, and $f^{tall}(w)(dave) >_{tall} f^{tall}(w)(phil)$. Thus, $\langle w, \delta_{tall}\rangle$ survives the update with with $^{\cap}$more(tall)(dave,phil), and so $\langle w, \delta_{tall}\rangle \in \sigma'$. Therefore, δ_{tall} is a degree available at σ'.

We conclude, then, that when the static semantics of the comparative offered in Kennedy 1997 is transported into our dynamic framework, we correctly predict the unique context change potential of sentences containing the comparative construction.

2.8 The Logic of Vague Predicates

When a linguist or philosopher says that he aims to "elucidate the logic of vague concepts", he typically means by this that he seeks to provide a semantics for vague predicates and to then deduce facts regarding the set of valid sentences his semantics determines. We have already introduced a semantics for vague predicates, as well as a property of "truth" applicable to sentences in a language containing vague predicates. We are therefore in a position to join in on all the fun of "elucidating the logic" of such a language. In Section 2.8.2, we will investigate both the set of validities our semantics determines and the notion of consequence most congenial to it.

First, however, we will explore further the concept of truth introduced in Definition 6.2. Theorists have often discussed the question of how the truth value of a complex sentence in a language containing vague predicates is derived from the truth value of its component sentences, especially in situations where a component sentence involves the application of a vague predicate P to one of its borderline cases. In Section 2.8.1, we will see how our semantic theory answers this question.

2.8.1 Sentence Connectives and the Projection of Indeterminacy

In this subsection, we describe how the truth value of a complex sentence is related to those of its constituent sentences [15].

Negation: The case of negation is easy.

ϕ	
T	**F**
F	**T**
I	**I**

Let σ be an information state composed of sequences s. If ϕ is T at σ, then it is supported by σ, hence $\neg\phi$ is rejected by σ and so is F at σ. If ϕ is F at σ, then it is rejected by σ, hence $\neg\phi$ will be supported by σ and so is T at σ. If ϕ is I at σ, then some sequences $s \in \sigma$ model ϕ and others do not model ϕ. Therefore, some sequences $s \in \sigma$ do not model $\neg\phi$ and others model $\neg\phi$. Therefore, $\neg\phi$ is I at σ.

Conjunction: The case of conjunction is more interesting.

[15] A note to readers familiar with the standard works in dynamic semantics: we do not here use any special dynamic interpretation of the logical connectives $\neg$, $\wedge$, $\vee$ or $\exists$. Instead, complex sentences are evaluated at sequences s within a state σ according to the usual recursive definition of the relation $\models$. Note, however, that since our semantics does not make use of "dynamic binding," the usual dynamic semantics for $\neg$, $\wedge$ and $\vee$ would be equivalent to what we use here.

$\phi \backslash \psi$	**T**	**F**	**I**
T	**T**	**F**	**I**
F	**F**	**F**	**F**
I	**I**	**F**	**F** / **I**

First, if ϕ is F at σ, then σ rejects ϕ. Therefore, $\phi \wedge \psi$ must be rejected by σ, and so the conjunction is F at σ.

Now, if ϕ is T at σ, then there are two cases to consider: either ψ is T at σ or ψ is I at σ. If ψ is T at σ, then σ supports both ϕ and ψ. Therefore, the conjunction $\phi \wedge \psi$ is supported by σ, and so is T at σ. If ψ is I, then some sequences $s \in \sigma$ model ψ and others do not ψ. Therefore, since ϕ is supported by σ, there are sequences $s \in \sigma$ which model $\phi \wedge \psi$ and ones which do not model $\phi \wedge \psi$. Thus, $\phi \wedge \psi$ is I at σ.

Finally, if ϕ is I at σ, there is only one case left to consider: ψ is also I at σ. At this point, however, we see that the truth functionality of conjunction breaks down [16]. The value of $\phi \wedge \psi$ depends on whether in σ the sequences which model ϕ overlap with the sequences which model ψ. Suppose that they do not overlap. Then there is no sequence $s \in \sigma$ which models both ϕ and ψ. Thus, every sequence $s \in \sigma$ models $\neg(\phi \wedge \psi)$, and so $\phi \wedge \psi$ is F at σ. Now suppose that they do overlap. Thus, there are sequences $s \in \sigma$ which model $\phi \wedge \psi$. However, since ϕ is I at σ, there must also be sequences $s \in \sigma$ which do not model ϕ and so do not model $\phi \wedge \psi$. Therefore, $\phi \wedge \psi$ is I at σ.

Let us quickly note some positive consequences of this non-truth-functionality which our semantics predicts for sentence connectives. It is often argued that a failing of the "fuzzy logic" analysis of vague predicates is that it predicts *all* conjunctions to bear an indeterminate truth value if both their conjuncts are assigned an indeterminate truth value. For some conjunctions, this is indeed the case. For example, accepting that SUV's and Winnebagos are borderline cases for the predicate "truck," the sentence "An SUV is a truck and a Winnebago is a truck" does strike one as being "uncertain" or indeterminately true. However, some speakers have the intuition that the sentence "An SUV is a truck and an SUV is not a truck" is a straightforward contradiction, and so should be assigned the value F [17]. Unlike the fuzzy logic analysis, our semantics correctly captures the difference between these two conjunctions. There are quite obviously states in which the sentence "An SUV is a truck and a Winnebago is a truck" is I, but there are no states in which the sentence "An SUV is a truck and an SUV

[16] This lack of truth functionality is a distinguishing feature of another semantic theory of vague predicates: supervaluationism. The connections between these ideas and supervaluationism will be explored in Section 2.8.3.

[17] The often discussed contrary intuition that this sentence can be intelligibly asserted will be addressed in Section 2.8.2.

is not a truck" is I. In all states this sentence is F, since in any state, due to the sentence's being a classical contradiction, every sequence within the state will fail to model it. Similarly, if it is, in some sense, a conceptual necessity that speakers employing the predicates "pink" and "red" know that no object can be both entirely pink and entirely red, then every state will be one which rejects the sentence "That shirt is entirely pink and that shirt is entirely red," even if the shirt in question is a borderline case for both "entirely pink" and "entirely red." This result cannot be reproduced by "fuzzy logic" analyses or any other semantic frameworks in which sentence connectives within a language containing vague predicates are assigned truth functions as their semantic value.

To use the standard jargon, our semantic analysis grabs hold of the "penumbral connections" stretching out between vague predicates. This point will be returned to in the comparison in Section 2.8.3 of our approach to the supervaluational analysis of vague language, another account which famously succeeds in preserving the penumbra.

Disjunction: For disjunction we may adopt the standard definition $\phi \vee \psi =_{df} \neg(\neg\phi \wedge \neg\psi)$. Although the following truth table follows immediately, a more extended explanation of how disjunction operates in this system will be illuminating and important for later discussion.

ϕ \ ψ	**T**	**F**	**I**
T	**T**	**T**	**T**
F	**T**	**F**	**I**
I	**T**	**I**	**T** / **I**

First if ϕ is T at σ, then σ supports ϕ. Therefore, $\phi \vee \psi$ is supported by σ, and so the disjunction is T at σ.

Now, if ϕ is F at σ, there are two cases to consider: either ψ is F at σ or ψ is I at σ. If ψ is F at σ, then σ rejects both ϕ and ψ, and so the disjunction $\phi \vee \psi$ will be rejected by σ, making it F at σ. If ψ is I, then some sequences $s \in \sigma$ model ψ and others do not model ψ. Therefore, since ϕ is rejected by σ, there are some sequences $s \in \sigma$ which model $\phi \vee \psi$ and some which do not model $\phi \vee \psi$. Thus, $\phi \vee \psi$ is I at σ.

Finally, if ϕ is I at σ, only one case need be considered: ψ is also I at σ. Again, at this point, the truth functionality of the connective breaks down. The value of $\phi \vee \psi$ depends upon whether in σ the sequences which do not model ϕ overlap with those sequences which do not model ψ. If those sequences do not overlap, then no sequence fails to model both ϕ and ψ and so every sequence in σ models one of the two. Thus every sequence $s \in \sigma$ models $\phi \vee \psi$, and so $\phi \vee \psi$ is T at σ. Now suppose that they do overlap. Then there are sequences which do not model $\phi \vee \psi$. Since ϕ is I, there are sequences $s \in \sigma$ which model ϕ and so model $\phi \vee \psi$. Thus, $\phi \vee \psi$ is I at σ.

We will quickly note in passing that the non-truth-functionality of disjunction which our semantics derives also allows the capture of certain penumbral connections. For example, although the sentence "An SUV is a truck or a Winnebago is a truck" strikes one as indeterminate because each of its constituent sentences is indeterminate, many speakers have the intuition that "An SUV is a truck or an SUV is not a truck" is a straightforward tautology and so should be assigned the value T. The reader may easily check that in any state the sentence "An SUV is a truck or an SUV is not a truck" is T, but that in some states "An SUV is a truck or a Winnebago is a truck" is I. Again, a fuzzy logic approach cannot reproduce this result.

Material Implication: Finally, we will include for completeness the following truth table for material implication. Once more, we simply adopt the standard definition $\phi \to \psi =_{df} \neg\phi \vee \psi$. Given our semantics for negation and disjunction, the following is an easy consequence.

$\phi \setminus \psi$	**T**	**F**	**I**
T	**T**	**F**	**I**
F	**T**	**T**	**T**
I	**T**	**I**	**T** / **I**

2.8.2 Validity and Consequence

In this subsection, we lightly explore the notions of "validity" and "consequence" which spring immediately from our definition of "truth."

In classical logic, a sentence is valid if it is true in all models. Thus, it would look suspicious if we did not adopt the following as the definition of "validity" for a language containing vague predicates.

Definition 8.1 A sentence is **valid** if it is T at all information states. A sentence is **classically valid** if it is true in all classical models.

Happily, in nearly all systems of dynamic semantics, a sentence is said to be valid if it is supported by every information state. Our concept of validity, then, coincides perfectly with that standardly adopted in dynamic frameworks. Therefore, we may, if we like, import over to our logic whatever general results are known concerning validity in systems of dynamic semantics.

No such results are presently known to the author. However, we may establish that, due to the "classicality" of the sequences composing our information states, a formula in a language containing vague predicates is valid if and only if it is classically valid.

Proposition 8.2 Let P be a vague predicate, O^P be its associated set of degrees, and ϕ be a sentence containing P. ϕ is T at every information state if and only if ϕ is classically valid.

PROOF: Suppose ϕ is a classical validity and $\sigma \subseteq W \times O^P$ is an arbitrary information state. Now, clearly, for every pair $\langle w, \delta \rangle \in \sigma$, there is some classical model $\mathcal{M}$ such that for all sentences ψ, $w, \delta \models \psi$ iff $\mathcal{M} \models \psi$ [18]. Therefore, ϕ is modelled by every pair $\langle w, \delta \rangle \in \sigma$, and so ϕ is T at σ. Since σ was arbitrary, ϕ is T at every information state.

Now suppose that ϕ is T at every information state. Then ϕ is T at $W \times O^P$. Clearly, for each classical model $\mathcal{M}$, there is some pair $\langle w, \delta \rangle$ such that for all sentences ψ, $w, \delta \models \psi$ iff $\mathcal{M} \models \psi$. Therefore, since ϕ is modelled by every pair $\langle w, \delta \rangle$, it is modelled by every classical model $\mathcal{M}$. So, ϕ is a classical validity.

Having settled on a notion of validity, the next task is to state what it means for a formula ψ to be a consequence of a formula ϕ. Again, given the definition of "consequence" in classical logic, we would appear goofy to offer anything but the following.

Definition 8.3 A sentence ψ is a **consequence of** a sentence ϕ if ψ is T at every information state at which ϕ is T. A sentence ψ is a **classical consequence of** a sentence ϕ if ψ is true in every classical model in which ϕ is true.

Although this notion of "consequence" is highly intuitive, it does differ from that which is most often adopted within systems of dynamic semantics. Typically, the definition of consequence in a dynamic framework must be made slightly more complicated, since these systems usually adopt a special semantics for quantifiers in order to model the natural behavior of anaphors and indefinites. In our study, however, we have not concerned ourselves with the vagaries of anaphora, and so our semantics can leave them out of the picture. Doing so allows us to adopt this simpler and more intuitive notion of consequence.

Not only is the proffered definition of "consequence" highly intuitive, it once again adheres perfectly to the classical notion.

Proposition 8.4 Let P be a vague predicate, O^P be its associated set of degrees, and ϕ, ψ be sentences either of which contains P. ψ is a consequence of

[18] After all, as noted in footnote 6, our language of vague predicates includes no logical symbols or other syntax besides that available in standard first order logic, and the logical constants of the language are interpreted at a sequences within a state exactly as they are within classical models.

ϕ iff ψ is a classical consequence of ϕ.

PROOF: Suppose that ψ is a consequence of ϕ. Let $\mathcal{M}$ be a classical model such that $\mathcal{M} \models \phi$. Now, there exists some $\langle w, \delta \rangle \in W \times O^P$ such that for all sentences χ, $\mathcal{M} \models \chi$ iff $w, \delta \models \chi$. Consider the state $\sigma = \{\langle w, \delta \rangle\}$. Clearly, ϕ is T at σ, and so by assumption, ψ is T at σ. Thus, $w, \delta \models \psi$ and so $\mathcal{M} \models \psi$. Since $\mathcal{M}$ was arbitrary, for any classical model $\mathcal{M}$, if ϕ is true in $\mathcal{M}$, then ψ is true in $\mathcal{M}$, and so ψ is a classical consequence of ϕ.

Now suppose that ψ is a classical consequence of ϕ, and suppose that ϕ is T at σ. For every $\langle w, \delta \rangle \in \sigma$, there exists a classical model $\mathcal{M}$ such that for all sentences χ, $\mathcal{M} \models \chi$ iff $w, \delta \models \chi$. Thus, every such $\mathcal{M}$ models ϕ and so models ψ. Therefore, every $\langle w, \delta \rangle \in \sigma$ models ψ, and so ψ is T at σ. Since σ was arbitrary, ψ is a consequence of ϕ.

Propositions 8.2 and 8.4 allow us to state the following neat result.

Proposition 8.5 (Deduction Theorem for Vague Language) Let P be a vague predicate, and ϕ, ψ be sentences either of which contain P. ψ is a consequence of ϕ iff $\phi \to \psi$ is valid.

PROOF: This follows trivially from Propositions 8.2, 8.4, and the Deduction Theorem for classical logic.

These results show that the logic of vague predicates which follows most naturally from our definition of truth in Section 2.6 is simply the logic which would result from assigning all vague expressions a classical semantics [19]. It may be argued, then, that our semantic analysis of vague predicates explains why, despite all the gnashing about by philosophers over the sorites paradox, all the principles of classical logic appear to be true of vague predicates as well as precise ones. That is, in response to the sorites paradox, philosophers sometimes recommend our abandoning certain principles of classical logic when our language employs vague expressions. As discussed in Cargile 1969 and Dummett 1975, however, such strategies always run into the snag that the principles they recommend we abandon still strike one as obviously sound for vague language when considered individually. If our semantic analysis of vagueness is on the right track, it would explain why a "resolution" to the sorites cannot be found in our abandoning certain strategies of reasoning. Instead, we must look to the

[19] The reader may note that a consequence of this *strict* classicality for our resulting logic is that certain intuitively valid inferences involving vague adjectives are not permitted within our system. Frank Veltman (p.c.) has mentioned the following example: all A's are B's; some big B's are small A's; therefore, all big A's are big B's. The failure of our semantics to capture the validity of this inference is due to its insufficiently modeling the way in which the standards for vague predicates are related to one another and to the properties of their implicit comparison classes. Although this is an important topic, one which dynamic analyses of vagueness have yet to fully address, it is – to use the cliché – beyond the scope of our present study.

possibility of denying one of the premises within the sorites argument. Just such a strategy will be promoted in Chapter 4.

On the other hand, this "classicality" of our logic for vague predicates may dissuade some from embracing the semantic analysis which justifies it. Some philosophers and linguists have argued that any proper semantic theory of vagueness will have to yield a non-classical logic of vague predicates, one in which there is a different set of validities and consequences from classical logic. Their conviction usually centers around the questioned validity of law of excluded middle (h.f., "LEM") and, as was just mentioned, the sorites paradox. Since we will discuss the sorites in Chapter 4, let us focus here on the first issue: the validity of LEM for vague language.

To say that the LEM is "invalid" for a language containing vague predicates means that there are sentences ϕ in such a language for which $\phi \vee \neg\phi$ is not necessarily true. Clearly, if one accepts the semantics we propose, then LEM will not be invalid for languages containing vague predicates; in every state, every sentence of the form $\phi \vee \neg\phi$ is T. However, some would maintain that this is an error of our theory.

The reason typically offered in support of the invalidity of LEM within a vague language is one we might call the "idiom's challenge". In everyday speech, when a borderline case α for some predicate P is encountered, this is often communicated by uttering the sentence "$P\alpha \wedge \neg P\alpha$" or the sentence "$\neg(P\alpha \vee \neg P\alpha)$". For example, we sometimes say "He's tall, and, then again, he isn't tall" or "He's neither really tall nor (really) not tall" when someone is borderline tall. Similarly, it has been claimed (Dummett 1975) that if one wishes to assert that a certain predicate P does *not* admit of borderline cases, one will often say "$\forall x(P(x) \vee \neg P(x))$". One might, for example, reprimand that "A person is either a doctor or he isn't (for crying out loud!)." Certainly, if LEM were invalid for vague language—in particular, if its invalidity were witnessed precisely by sentences $P(\alpha)$ in which α is a borderline case for P—this would make sense of such behavior.

However, other explanations are preferable. As is often the case with data from natural language, intonation is everything. Even if α is a borderline case for P, one still cannot say either $P(\alpha) \wedge \neg P(\alpha)$ or $\neg(P(\alpha) \vee \neg P(\alpha))$ with natural speed in a normal, declarative tone. For example, the sentence "Dave is tall, and he isn't tall," when spoken in the same manner as "Dave is Jewish, and he's from New Jersey" simply sounds bizarre, and one would want to label it an outright contradiction. Instead, to attain felicity, one must pronounce this sentence with a very slow, uncertain intonation, and give a considerable amount of pause between the first and second conjunct. Furthermore, even when using such a manner of speech, there are sentences of the logical forms $P(\alpha) \wedge \neg P(\alpha)$ and $\neg(P(\alpha) \vee \neg P(\alpha))$ which remain infelicitous. To my ear, there is no way to pronounce the sentence "An SUV is both a truck and not a truck" which communicates that an SUV is a borderline case for the predicate "truck." Perhaps due to the presence of "both," this sentence strikes me as a blatant contradiction, no matter the context or style of delivery. A semantics for vague predicates which withheld the LEM from sentences in which a predicate is applied to one of its borderline cases would not be able to explain these facts. Indeed, such a theory would seem to predict that, as factual assertions, certain sentences of the logical form $\neg(P(\alpha) \vee \neg P(\alpha))$ should receive a normal, declarative intonation. Moreover, the existence of any restrictions on the surface syntactic form of these sentences is, from such a point of view, completely mysterious.

Thankfully, other explanations of these facts are available. One analysis, which we shall only sketch here, is that when one asserts "He's tall, and (then again) he isn't tall" or "He's neither really tall nor (really) not tall", the utterance may not be *true* as a literal assertion, but it can as a speech act communicate an inability to decide on the truth of the constituent sentences. From the point of view of one uttering these sentences, both of the constituent propositions are equally assertible. The two expressions differ, however, on the attitude towards the constituent propositions which the speaker communicates; the first idiom expresses the speaker's willingness to assert either, while the second idiom expresses the speaker's reluctance to assert either. We find exactly such uses of these sentences in, for example, the utterance of "I want to go and, then again, I don't want to go" by a speaker who is vacillating before a decision to leave. Clearly, such a speaker is not a borderline case of "wanting to go". Everyone would agree that the speaker does not (wide scope) want to go; otherwise, he would not be vacillating. Thus, in this context, the sentence clearly communicates the speaker's indecision, not the fact that the speaker is a borderline case for some vague predicate. Likewise, when one asserts in a slow, cautious intonation that "Dave is tall, and (then again) Dave is not tall", one is expressing, giving voice to, acting out their indecision to assert either conjunct [20]. One is not asserting of Dave that he has two complementary properties, which is why this sentence does not get the normal, declarative intonation.

Similarly, if a speaker asserts $P \vee \neg P$, his assertion may be an empty tautology, but it can be used rhetorically to remind his listener of this universally applicable logical principle. If one denies both the claims P and $\neg P$, and it is understood that one is not engaging in the "indecisive" speech act just described, then one is open to a charge of logical inconsistency. This charge may be levied through an assertion of the LEM. That is, contrary to the claims of Dummett 1975, assertions of the form $P \vee \neg P$ seem *not* to communicate that P is a precise predicate, but that the listener appears to the speaker to be contradicting themselves in denying both of the disjuncts. For example, "one is either a civilian or not a civilian" fits most naturally into a scenario where it is hurled at a retired military man who attempts both to obtain some privilege granted only to acting military personnel, and to deny his having to report to military duty. It is, in fact, beyond my ability to imagine a scenario in which $P \vee \neg P$ is clearly being used to communicate the precision of the predicate P. Furthermore, the very idea that an assertion of the LEM communicates that a predicate is precise is undermined by the fact that although virtually *no* pred-

[20]Another analysis of this speech act, perhaps equally plausible, is that one is acting out their *freedom* to assert either sentence. Whichever analysis one prefers, the greater point remains.

icates in natural language are precise, the idiom itself has a wide applicability. That is, if the assertion "a person is either a P or he isn't" may be used to mean that "P" does not admit of borderline cases, then that use must be incredibly rare, since nearly all predicates in natural language (including "doctor" and "civilian") admit of borderline cases, but assertions of this form appear to memory as being often sincerely made.

In short, asserting a sentence of the logical form $P \wedge \neg P$, $\neg(P \vee \neg P)$, or $P \vee \neg P$ may have various rhetorical effects, but one's doing so is not necessarily indicative of either the semantic value of those expressions or the general validity of LEM in vague language. We conclude, then, that our semantics' justification of the application of classical logic to vague predicates is a strong point in its favor.

2.8.3 Connections with Supervaluationism

The reader has perhaps already noticed a similarity between our dynamic semantics for standard-sensitive context dependent predicates and the supervaluational analysis of vague predicates, first suggested by Henryk Mehlberg (Mehlberg 1958) and given its mature form by Kit Fine (Fine 1975) [21]. To recall, this analysis holds that a sentence ϕ containing a vague predicate P is true iff the sentence is "supertrue", iff it is T in all ways of making P "totally precise." A language $\mathcal{L}$ containing vague predicates is interpreted over a relational structure, $\langle W, \leq \rangle$ in which W is a set of points associated with three valued interpretations of $\mathcal{L}$ and $\leq$ is the relation of "more precise." Within the set W there are always a number of points $W_{classical} \subset W$ associated with two valued interpretations of $\mathcal{L}$, interpretations of $\mathcal{L}$ in which sentences are assigned only the values T or F. Due to the properties stipulated to hold of $\leq$, these classical points are always limit points of the chain induced by $\leq$. Roughly, a sentence of $\mathcal{L}$ is said to be "supertrue" in a structure $\langle W, \leq \rangle$ if it is T at every $w \in W_{classical}$. Truth for sentences in a vague language is then equated with supertruth in a relational structure.

The supervaluational theory of vague language was developed mainly to provide an analysis for vague predicates which accurately captures the penumbral connections holding between them. For example, we noted earlier that although the sentences "Dave's shirt is entirely red" and "Dave's shirt is entirely pink" may in some cases both be of indeterminate truth value, the sentence "Dave's shirt is entirely red and Dave's shirt is entirely pink" seems in all cases to be false. Similarly, whatever the truth value of the sentence "An SUV is a truck," the sentence "An SUV is a truck or an SUV is not a truck" sounds to be a tautology. The supervaluational analysis captures these facts via its identification of "truth" with "supertruth" and the various constraints placed upon the relation $\leq$. These constraints are understood as limits on the way in which the meaning of a vague predicate may be "extended" and made more precise. One

[21] Although "supervaluations" were first introduced in the work of van Fraassen (van Fraassen 1969), the core ideas of the supervaluational analysis of vague language appear first in Mehlberg's work, though he himself does not use the term "supervaluation."

such constraint would be that, in whatever way "red" is made more precise, it cannot overlap in its extension with "pink." Formally, this constraint is stated as the condition that for all objects α, if red(α) is T at a point w, then for all points $w \leq w'$, red(α) is T at w' and pink(α) is not T at w'. The result of this constraint is that at all $w \in W_{classical}$, $w \models \forall x(red(x) \rightarrow \neg pink(x))$, and therefore $w \models \forall x \neg(red(x) \wedge pink(x))$. Thus, the sentence "Dave's shirt is entirely red and Dave's shirt is entirely pink" is not supertrue, and so is not true, since it is false at every $w \in W_{classical}$. Similarly, since every point $w \in W_{classical}$ is associated with a classical, two-valued model, at every $w \in W_{classical}$ any sentence of the form $\phi \vee \neg\phi$ is T and any sentence of the form $\phi \wedge \neg\phi$ is F. The supervaluational analysis therefore accurately predicts that the sentence "An SUV is a truck and an SUV is not a truck" is a contradiction, while "An SUV is a truck or an SUV is not a truck" is a tautology.

The reader may now be tempted to make a very simple equation between the supervaluational analysis and the dynamic semantics proposed here: the points $W_{classical}$ are simply the pairs constituting the state σ, and supertruth is simply support. Although there is certainly something to this analogy, and the two theories do bear a strong resemblance in that they both justify the application of classical logic to vague language, I would like to wave my hands at some finer ways in which they differ.

First, and most importantly, the two theories differ in their motivation. The supervaluational analysis is an attempt to develop a semantics for vague language that adequately captures the existence of penumbral connections. Our dynamic semantics for standard-sensitive predicates is most concerned with the question of why that form of context dependency and the possibility of having borderline cases are so closely linked. Interestingly, our dynamic semantics also provides a means for handling penumbral connections, which at least puts it on equal footing with the supervaluational theory. However, the supervaluational theory does not clearly account for the connection between standard-sensitivity and borderline cases. In fact, on the supervaluational analysis, there is no reason why predicates with borderline cases shouldn't have the same, "fuzzy" extension in every context. Our dynamic semantics says a little more about the nature of borderline cases than the supervaluational analysis, and can therefore build the bridge between vagueness and context dependency. Furthermore, it will be shown in Chapters 3 and 4 how this basic dynamic analysis may be extended to offer unique analyses of higher order vagueness and the sorties paradox.

These two theories also differ in their understanding of the "penumbral connections" holding between vague predicates. The supervaluationist analysis claims that penumbral connections are, in part, a kind of superstructural constraint on the way in which vague predicates can be made more precise. The dynamic account simply holds that they are facts the speakers know. It is sometimes claimed by philosophers that there are some things a speaker *must* believe in order for him to use certain predicates correctly. If a speaker of English must know that something entirely red cannot be entirely pink, then any information state for a speaker of English will support the formula $\forall x \neg(red(x) \wedge pink(x))$. Thus, any information state will reject an utterance of "Dave's shirt is entirely red and entirely pink", and the formula is judged F. From our perspective, then, penumbral connections aren't so much conditions on making language precise as they are merely facts which speakers of a language must know.

A final distinction worth pointing out is that the property of truth in our dynamic semantics is not perfectly analogous to the notion of "supertruth" in the supervaluational analysis. The latter account bases the truth of an expression on imagined ideal states of the language. The dynamic account bases the truth of an expression on the knowledge the listener has of the context. Truth in the dynamic account is not in any sense "truth from above," "truth in an imagined ideal". Our dynamic analysis uses nothing more than the concept of truth most natural to dynamic semantics: it's true if your information entails it. It is an interesting fact that this concept behaves quite similarly to the notion of "supertruth".

Let us conclude by noting in passing that despite these rather "philosophical" differences, the connection between the supervaluational semantics of vague predicates and our dynamic semantics is still largely unexplored and not fully understood. One may, at this point, if one likes, equate in one's mind the formalism of our dynamic semantics and that of the supervaluational analysis. In the following chapters, however, this equation will only become more strained.

2.9 The Semantics of "Definitely"

For reasons that will become vividly clear at the start of the next chapter, most works proposing a semantic analysis of vague predicates at some point discuss the expressions "it is definitely the case that" and "it is indefinite whether." Nearly all theories represent these expressions within their formal metalanguages as sentence operators. Tradition will not be bucked on this point, and so in this section we will introduce such operators into our language and discuss their semantics.

The first subsection is intended to "whet" our imaginations concerning the semantics of these operators, before we jump into the full, proffered analysis in Section 2.9.2.

2.9.1 A First, Crude Pass

The operator "D" is to be read as "it is definitely the case that." We add it to our language with the following syntactic clause.

Definition 9.1 If ϕ is a formula of our language, then $D\phi$ is a formula of our language.

With the introduction of D into our language, we may now also introduce the operator I, read as "it is indefinite whether" and given the following, standard definition.

Definition 9.2 $I\phi =_{df} \neg D\phi \wedge \neg D\neg\phi$

Now, the intuitive meanings of D and I place certain restrictions on the semantic definition we offer for D. ϕ should be T if $D\phi$ is T, and if ϕ is not T, then $D\phi$ should not be T either. Given our definition of "truth," we might offer the following semantics for D as a first pass.

Definition 9.3 Let $\sigma \subseteq W \times O^P$ be an information state. $[\![D\phi]\!](\sigma) = \sigma$ if $[\![\phi]\!](\sigma) = \sigma$. Otherwise, $[\![D\phi]\!](\sigma) = \emptyset$.

That this definition is sufficient for capturing the intuitive relationship between the truth of $D\phi$ and ϕ may be seen via the following proposition.

Proposition 9.4 Let $\sigma \subseteq W \times O^P$ be an information state. $D\phi$ is T at σ iff ϕ is T at σ. $D\phi$ is F at σ otherwise.

PROOF: By Definition 6.1 and Definition 6.2, $D\phi$ is T at σ iff $[\![D\phi]\!](\sigma) = \sigma$, iff, by Definition 9.3, $[\![\phi]\!](\sigma) = \sigma$, iff, by Definition 6.1 and 6.2, ϕ is T at σ.

Now suppose that ϕ is not T at σ. By Definition 6.1 and 6.2, $[\![\phi]\!](\sigma) \neq \sigma$, and so by Definition 9.3, $[\![D\phi]\!](\sigma) = \emptyset$. Finally, by Definition 6.1 and 6.2, $D\phi$ is F at σ.

With this primitively dynamic semantics for D, however, we must alter slightly our present interpretation of the sentence connectives, adopting instead a primitively dynamic interpretation for them as well.

Definition 9.5 If ϕ is a sentence of the form $\neg\psi$, then $[\![\phi]\!](\sigma) = \{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\psi]\!](\sigma)\}$. If ϕ is a sentence of the form $\psi \wedge \chi$, then $[\![\phi]\!](\sigma) = [\![\chi]\!]([\![\psi]\!](\sigma))$.

Proposition 9.4 clearly entails that ϕ is T at σ if $D\phi$ is, and $D\phi$ is F at σ if ϕ is F. Another pleasant result of this semantics is that $I\phi$ is T at a state σ if and only if ϕ is I at σ.

Proposition 9.6 $I\phi$ is T at σ if and only if ϕ is I at σ.

PROOF: Suppose that ϕ is I at σ. Then, there are sequences $s \in \sigma$ which model ϕ and ones which do not. Thus, $[\![\phi]\!](\sigma) \neq \sigma$ and $[\![\neg\phi]\!](\sigma) \neq \sigma$, and so $D\phi$ is F at σ as well as $D\neg\phi$. By Definition 9.5, then, $\neg D\phi$ and $\neg D\neg\phi$ are both T at σ, and so by Definitions 9.2 and 9.5, $I\phi$ is T at σ.

Now suppose that ϕ is not I at σ. Thus, ϕ is either T or F at σ. Suppose that ϕ is T at σ. Then, by Definition 9.3 and Definition 9.5, $[\![\neg D\phi]\!](\sigma) = \emptyset$. Moreover, by Definition 9.3 and 9.5, $[\![\neg D\neg\phi]\!](\emptyset) = \emptyset$. Thus, by Definition 9.2 and 9.5, $[\![I\phi]\!](\sigma) = \emptyset$ and so $I\phi$ is not T at σ. Suppose that ϕ is F at σ. Then, by Definition 9.3 and 9.5, $[\![\neg D\phi]\!](\sigma) = \sigma$. However, by Definition 9.3 and 9.5, $[\![\neg D\neg\phi]\!](\sigma) = \emptyset$, which entails by Defintion 9.2 and 9.5, that $[\![I\phi]\!](\sigma) = \emptyset$. Thus,

$I\phi$ is not T at σ.

Our semantics for D thus captures an intuitive relationship between the truth value of a sentence ϕ and the truth values of the sentences $D\phi$ and $I\phi$. Moreover, it also correctly relates the meaning of "it is definitely the case that ϕ" to a speaker's knowing that ϕ. For example, if one says that Dave is "definitely tall," one means that there is no possible doubt concerning Dave's tallness. Formally, we would represent this situation as an information state σ at which tall(dave) is supported, and so our semantics stands as a plausible analysis of these natural language expressions.

Happily, the addition of D and I to our language does not seriously affect the nice results obtained in the last section. For example, we have not lost the deduction theorem for our language.

Proposition 9.7 Let ϕ, ψ be any two sentences of our extended language. ψ is a consequence of ϕ iff $\phi \to \psi$ is valid.

PROOF: Recall that $\phi \to \psi$ is defined in section 8.1 as $\neg(\neg\neg\phi \wedge \neg\psi)$, which by Definition 9.5 is equivalent to $\neg(\phi \wedge \neg\psi)$.

First, let us suppose that ψ is a consequence of ϕ, and let $\sigma \subseteq W \times O^P$ be an arbitrary information state. At σ, ϕ is either T, F or I. Now, suppose that ϕ is T at σ. Let us consider the value of $[\![\neg(\phi \wedge \neg\psi)]\!](\sigma) = \{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!]([\![\phi]\!](\sigma))\}$. Since, by assumption, $[\![\phi]\!](\sigma) = \sigma$, $[\![\psi]\!](\sigma) = \sigma$. Therefore, $\{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!]([\![\phi]\!](\sigma))\} = \{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin \emptyset\} = \sigma$. Thus, $\phi \to \psi$ is T at σ. Now, suppose that ϕ is F. Then, $[\![\phi]\!](\sigma) = \emptyset$, and so $\{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!]([\![\phi]\!](\sigma))\} = \{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin \emptyset\} = \sigma$. Thus, $\phi \to \psi$ is T at σ. Finally, suppose that ϕ is I. Then $[\![\phi]\!](\sigma) = \sigma' \neq \emptyset$. Now, since $[\![\phi]\!](\sigma') = \sigma'$, then by assumption $[\![\psi]\!](\sigma') = \sigma'$. Therefore, $\{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!]([\![\phi]\!](\sigma))\} = \{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!](\sigma')\} = \{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin \emptyset\} = \sigma$. Thus, $\phi \to \psi$ is T at σ. Since σ was arbitrary, $\phi \to \psi$ is valid.

Now, let us suppose that $\phi \to \psi$ is valid, and let $\sigma \subseteq W \times O^P$ be a state at which ϕ is T. Now, since $\phi \to \psi$ is valid, then $\{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!]([\![\phi]\!](\sigma))\} = \sigma$. Moreover, by definition of σ, $\{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!](\sigma)\} = \sigma$. Thus, $\sigma \cap [\![\neg\psi]\!](\sigma) = \emptyset$. However, by the nature of updates, $[\![\neg\psi]\!](\sigma) \subseteq \sigma$. Therefore, $[\![\neg\psi]\!](\sigma) = \emptyset$, and so $\psi(\sigma) = \sigma$. Thus, ψ is T at σ. Since σ was arbitrary, ψ is a consequence of σ.

Given the tight connection between the truth of ϕ and that of $D\phi$ demonstrated in Proposition 9.4, the relation of consequence between sentences ϕ, ψ within our extended language can also be characterized by the formula $D\phi \to \psi$.

Proposition 9.8 (Quasi-Deduction Theorem for Vague Language) Let P be a vague predicate, and let ϕ, ψ be sentences either of which contain P. ψ is a consequence of ϕ iff $D\phi \to \psi$ is valid.

PROOF: First, suppose that ψ is a consequence of ϕ, and let $\sigma \subseteq W \times O^P$ be an arbitrary information state. ϕ is either supported or not supported by σ. Suppose that ϕ is supported by σ. Consider the value of $[\![D\phi \to \psi]\!](\sigma) = [\![\neg(D\phi \wedge \neg\psi)]\!](\sigma) = \{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!]([\![D\phi]\!](\sigma))\}$. Since ϕ is supported by σ, then $[\![D\phi]\!](\sigma) = \sigma$, and so $\{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!]([\![D\phi]\!](\sigma))\} = \{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!](\sigma)\}$. Moreover, since ψ is a consequence of ϕ, then $[\![\psi]\!](\sigma) = \sigma$. Thus, $\{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!](\sigma)\} = \{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin \emptyset\} = \sigma$. Therefore, $D\phi \to \psi$ is T at σ. Now, suppose that that ϕ is not supported by σ. Thus, $\{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!]([\![D\phi]\!](\sigma))\} = \{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!](\emptyset)\} = \{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin \emptyset\} = \sigma$. Therefore, $D\phi \to \psi$ is T at σ. Since σ was arbitrary, $D\phi \to \psi$ is valid.

Now, suppose that $D\phi \to \psi$ is valid, and let $\sigma \subseteq W \times O^P$ be an information state at which ϕ is T. Thus, $[\![\phi]\!](\sigma) = \sigma$, and by assumption, $\{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!]([\![D\phi]\!](\sigma))\} = \sigma$. Therefore, $\{\langle w, \delta\rangle \in \sigma : \langle w, \delta\rangle \notin [\![\neg\psi]\!](\sigma)\} = \sigma$. By reasoning similar to that in the proof of Proposition 9.7, we conclude that $[\![\psi]\!](\sigma) = \sigma$, and so ψ is T at σ. Since σ was arbitrary, ψ is a consequence of ϕ.

With this semantics for D and I, our extended language can express a limited range of statements about the information which speakers have at a context. We are immediately faced with a problem, however, if we wish our semantics for D to capture the widespread intuition that there exists a difference between something being "definitely" the case and it being "definitely definitely" the case. That is, we would like for there to be states in which $D\phi$ is T but not $DD\phi$. Proposition 9.4 already shows that this is out of the question. The source of this difficulty is that at any state, sentences of the form "$D\phi$" will have only either the value T or F. But, it seems that it could be indeterminate whether the proposition "it is definitely the case that ϕ" is true, and so it should be possible for $D\phi$ to be assigned the value I. An immediate nasty consequence of this limitation on the truth values of $D\phi$ is revealed in the following proposition.

Proposition 9.9 $II\phi$ cannot be T at any information state.

PROOF: Let $\sigma \subseteq W \times O^P$ be an information state. Now, at σ, the sentence ϕ is either T, F or I. Suppose that ϕ is T at σ. Then, by the reasoning laid out in the proof of Proposition 9.6, $[\![I\phi]\!](\sigma) = \emptyset$ and so $I\phi$ is F at σ. Now, since $I\phi$ is not I at σ, by Proposition 9.6, $II\phi$ cannot be T at σ. Suppose that ϕ is F at σ. Again, by the reasoning laid out in the proof of Proposition 9.6, $[\![I\phi]\!](\sigma) = \emptyset$ and so $I\phi$ is F at σ. Again, it follows from Proposition 9.6 that $II\phi$ cannot be T at σ. Finally, suppose that ϕ is I at σ. Then, by Proposition 9.6, $I\phi$ is T at σ, and so by Proposition 9.6, $II\phi$ cannot be T at σ.

In section 3.2.2, we will see that Proposition 9.9 reveals a critical failure of our present semantics for D and I. In the next subsection, we therefore introduce a more sophisticated semantics which resolves these tensions.

2.9.2 The Preferred Approach

The purpose of annoying the reader with the bitterly flawed semantics for D proposed in the previous subsection was to convince him that, due to our use in this thesis of a dynamic framework, a simpler approach to the semantics of D and I than what will be proposed in this section is not likely to be possible.

To build some momentum towards our proposal, let us first take note of the fact, mentioned in the previous subsection, than an assertion of the form $D\phi$ seems to express a fact about the speaker's knowledge, or confidence. One thing we all know about the actual world w_0 is that there are a number of conversations taking place within it, all with different contexts and knowledge states holding for the speakers. Each of these conversations and contexts could have, of course, been different. This suggests that the possible worlds within the sequences making up our knowledge states should contain as "constituent facts" the properties of any conversations being held "within" them, including the content of their contexts. A particular possible world, then, should be able to model or refute a statement such as "We all know (in this conversation) that Laura is fat" or "It is definitely the case (in that conversation) that John is tall." This, too, relates to the fact that each speaker within a conversation may have beliefs regarding the context which are less than accurate. If σ is the context for a conversation c at a world w, and all the worlds within sequences making up the knowledge state σ model the statement that the conversation c is one in which, say, Sam told a spicy joke, then this would intuitively represent that the speakers within c at world w believe that at some point in c Sam told a spicy joke. Moreover, uttering in a conversation "It is definitely the case that Sam told a spicy joke" seems to affect both one's knowledge of the world and *one's knowledge of the context*. One not only eliminates from their mind the possibility of Sam's not having told a spicy joke, but also the possibility of there being doubt over this fact.

Let us begin to bring these ideas into our dynamic semantics by introducing into our formalism the notion of a "conversation." A conversation c_i will be construed as a function from worlds to information states [22]. We represent that at world w the conversation c_i has the associated context σ by writing $c_i(w) = \sigma$. Now, instead of having the single operator D, we give to D a syntax much like the hybrid logic operator @. That is, we index D with the name of a conversation, creating a whole series of operators D_{c_1}, D_{c_2}, D_{c_3}, ..., one for each conversation c_i.

Let us now add this series of operators to our language with the following syntactic clause.

Definition 9.10 If ϕ is a formula of our language, then so is $D_{c_i}\phi$.

[22] In a way, then, conversations are represented as the "concepts" of information states. However, another way to characterize our formalism here is that conversations are identified rigidly across possible worlds, with the function c_i representing the way in which the information state of the conversation c_i changes from world to world.

For each "definitely operator" D_{c_i}, we also introduce an accompanying "indefinitely operator" I_{c_i}, defined as follows.

Definition 9.11 $I_{c_i}\phi =_{df} \neg D_{c_i}\phi \wedge \neg D_{c_i}\neg\phi$.

We now have all the tools necessary to provide an accurate dynamic semantics for the expressions "it is definitely the case that" and "it is indefinite whether." In this thesis, whenever our dynamic semantics employs the formalism of the "Stalnakerian Approach" introduced in Section 2.2, the operator D_{c_i} will be interpreted as follows.

Definition 9.12 Let P be a standard-sensitive context dependent predicate, O^P be its associated set of degrees, $\delta \in O^P$, and $\sigma \subset W \times O^P$ be an information state. If $\langle w, \delta\rangle \in W \times O^P$, then $\langle w, \delta\rangle \models D_{c_i}\phi$ iff ϕ is supported by the information state $c_i(w)$. The sentence $D_{c_i}\phi$ is given as its semantic value the function $[\![D_{c_i}\phi]\!]$ defined as follows.

$$[\![D_{c_i}\phi]\!] : \wp(W \times O^P) \rightarrow \wp(W \times O^P)$$
$$[\![D_{c_i}\phi]\!] : \sigma \mapsto \sigma \cap \{\langle w, \delta\rangle : \langle w, \delta\rangle \models D_{c_i}\phi\}$$

This semantics correctly captures the effect mentioned above of uttering a sentence "it is definitely the case that ϕ." Such an utterance removes from the knowledge state of the listener any sequence which does not model the sentence $D_{c_i}\phi$, that is, any sequence in which the context σ of the conversation c_i at the possible world of the sequence does not support ϕ. The result of making the utterance "it is definitely the case that ϕ", then, is that the listener's knowledge state becomes one in which he believes that the conversation contains the information that ϕ.

When a single, fixed conversation is the only one relevant to our interests, we will drop the subscript from the function c_i, and the operator D_{c_i} will be written simply as D.

Definition 9.13 Let c be the only conversation, P be a standard-sensitive context dependent predicate, O^P be its associated set of degrees, $\delta \in O^P$, and $\sigma \subseteq W \times O^P$ be an information state for the conversation c. If $\langle w, \delta\rangle \in W \times O^P$, then $\langle w, \delta\rangle \models D\phi$ iff ϕ is supported by the information state $c(w)$. The sentence $D\phi$ is given as its semantic value the function $[\![D\phi]\!]$ defined as follows.

$$[\![D\phi]\!] : \wp(W \times O^P) \rightarrow \wp(W \times O^P)$$
$$[\![D\phi]\!] : \sigma \mapsto \sigma \cap \{\langle w, \delta\rangle : \langle w, \delta\rangle \models D\phi\}$$

The reader may have observed that our semantics for D does not yet capture the intuition mentioned in the last subsection that ϕ should be true if $D\phi$ is true, and that $D\phi$ ought not to be true if ϕ isn't. We see now that this intuition rests on the assumption that speakers in a conversation have a certain form of positive introspection. This may be obtained through the following condition.

Condition 9.14 Let c be our fixed conversation and $\sigma \subseteq W \times O^P$ be an information state for c. There exists some $\langle w, \delta \rangle \in \sigma$ such that $c(w) = \sigma$.

Condition 9.14 is now enough for our semantics to yield the desired connection between the truth of a formula ϕ and the truth of the formulas $D\phi$ and $I\phi$.

Proposition 9.15 Let c be our fixed conversation and $\sigma \subseteq W \times O^P$ be an information state for c. If $D\phi$ is T at σ, then ϕ is T at σ.

PROOF: Suppose $D\phi$ is T at σ. Then $D\phi$ is supported by σ, and so $D\phi$ is modelled by all $\langle w, \delta \rangle \in \sigma$. By Condition 9.14, then, $D\phi$ is modelled by the $\langle w, \delta \rangle \in \sigma$ such that $c(w) = \sigma$. But then, σ must support ϕ and so ϕ is T at σ.

Proposition 9.16 Let c be our fixed conversation and $\sigma \subseteq W \times O^P$ be an information state for c. If ϕ is not T at σ, then $D\phi$ is not T at σ. In particular, if ϕ is I, then $D\phi$ cannot be T.

PROOF: Follows trivially from Proposition 9.15.

Proposition 9.17 Let c be our fixed conversation and $\sigma \subseteq W \times O^P$ be an information state for c. If $I\phi$ is T at σ, then ϕ is I at σ.

PROOF: Suppose $I\phi$ is T at σ. Then $I\phi$ is supported by σ, and so $I\phi$ is modelled by all $\langle w, \delta \rangle \in \sigma$. By Condition 9.14, then, $I\phi$ is modelled by the $\langle w, \delta \rangle \in \sigma$ such that $c(w) = \sigma$, and so σ supports neither ϕ nor $\neg\phi$. Thus, ϕ is I at σ.

As the reader may check, however, our semantics is not enough to establish the converse of the claims above. This is, in fact, a positive result. From the mere fact that a sentence is true at a state, it should not thereby follow that it is definitely true. Similarly, from the fact that a sentence is not definitely true, it should not follow that the sentence is not true, and although the truth of "it is indeterminate whether ϕ" should imply ϕ's being of indeterminate truth, ϕ might be of indeterminate truth within a context without the speakers agreeing on this fact.

Our analysis also predicts that the result of uttering the sentence "It is definitely the case that John is tall" should be a knowledge state that supports "John is tall." The dynamic semantics given in 9.12 and 9.13 requires that the result of uttering "It is definitely the case that John is tall" be a state in which every sequence within the state represents the conversation as being one in which it is known that John is tall. However, Condition 9.14 requires that every state σ contain some sequence in which the conversation is represented as having σ as its associated context. Therefore, if σ is a state at which it is not yet known that John is tall, an utterance of "it is definitely the case that John is tall" is predicted to trigger an accommodation process whereby σ is altered to include the information that John is tall, allowing for the update of "it is definitely the case that John is tall" to produce a state σ' in conformity with condition 9.14.

Certainly, our semantics for D and I do not perfectly capture the meanings of "it is definitely the case that" and "it is indeterminate whether," but our analysis here does overcome the problems of the first, crude pass made in the previous section [23]. The reader is invited to construct contexts in which $D\phi$ is assigned the value I, and ones in which $II\phi$ is T. Most importantly, however, our semantics correctly represents these statements as essentially communicating information regarding the speakers' knowledge of their own shared knowledge. Therefore, modulo various changes within our formalism, this will be the semantics for D and I adopted throughout this thesis.

[23]The reader may check, however, that under our new semantics, our language loses both the Deduction Theorem and the Quasi-Deduction Theorem. The invalid $DP(\alpha) \rightarrow P(\alpha)$ refutes both of these statements.

Chapter 3

Vague Predicates, Vague Selections

The Full Analysis

3.1 Chapter Overview

The first half of this thesis was largely preparation for the more focused work of its second half. The first chapter explained and justified the goals of the thesis; the second chapter proposed a simple dynamic semantics for vague predicates and demonstrated that it had a scattered bunch of desirable consequences. In this chapter, however, we begin using the analysis of Chapter 2 to treat in detail some of the deeper and more challenging questions surrounding vague language. Although this chapter touches a few different matters, lurking behind all that we do here is the problem of so-called "higher order vagueness."

We begin in Section 3.2 with a roll-call of problems for our basic analysis, one failing being that it does not capture a generalization we dub the "Law of Higher Order Vagueness." In a real sense, all the problems mentioned here may be reduced to that failure, which itself springs from our not having properly constrained the contexts associated with conversations in which vague language is used. The rest of Chapter 3 is basically a pursuit of the idea that a particular theory of the origins of vagueness may help us to constrain our information states in a way that will allow us to capture the targeted law. This theory would seek to reduce linguistic vagueness to "vagueness" of a certain non-linguistic character, and so in Section 3.3 we discuss the interest in and precedent for such a reduction.

Section 3.4 introduces the concept of a "vague selection," which we will later argue is the underlying source of linguistic vagueness, and some important features of these things are discussed. The link between vagueness and vague selections will require our taking the degree orderings of the last chapter *very* seriously, and so 3.5 is an attempt to convince the reader that we do not lapse into schizophrenia by doing so.

In Section 3.6, we attempt to communicate at a very informal level the perceived link between vague selections and linguistic vagueness. This link is vaguely described as the condition that every use of a vague predicate "depends upon" some vague selection from the degrees for that predicate. In Section 3.7 we design a formal representation of vague selections, and with it we precisely formulate as a constraint on information states the "dependency" relation informally stated to exist between vague selections and the use of vague predicates. This constraint is then shown in Section 3.8 to be all we need for our account to capture the Law of Higher Order Vagueness and so to overcome the difficulties raised in the next section.

3.2 Why This is Not Enough

Nothing in this world is without its flaws, and the deepest flaws let you see the farthest. This section will lay out some difficulties for our semantic analysis which will quickly spur us towards an interesting discussion of the metaphysics of vagueness. All the problems listed here have a common source: without the benefit of additional assumptions, our dynamic semantics predicts uses for vague predicates which simply do not exist.

3.2.1 Non-Existent Contexts

The word "fat" can never be used with a precise meaning, not ever. There may be conversations in which one's use of "fat" is clearer, or more precise than in other conversations. For example, if you and I have been speaking of fat persons for a while now, accepting and rejecting various individuals as "fat," and I were to say to you "Oh man, some fat guy just totally walked into the room," then you will have a much clearer idea of what to expect of the person eliciting my utterance than if we had been for the past hour entirely avoiding the subject of obesity. Nevertheless, philosophers and linguists have often voiced the intuition that a context in which "fat" is used with a fully precise meaning is simply unimaginable, except in cases where the speakers have thereby created a homophonous technical term. A more concrete way of stating this intuition is that in all plausible contexts a vague predicate such as "fat" must admit of *some* imaginable borderline cases.

This inability for vague predicates to be used precisely, which we might call the "ineluctability of vagueness," is often raised as a point against the supervaluational analysis of vague predicates. It is argued that the equation which that analysis makes between the truth of a sentence in a vague language and supertruth, truth for all ways of making the language fully precise, is straightforwardly negated by the ineluctability of vagueness. There are no ways of making the meaning of a vague predicate fully precise, and so the notion of "supertruth" is suspect, as well as the claim that the semantics of a vague langauge depends upon it.

Similar problems also arise for our dynamic semantics. Our theory of information states permits the existence of contexts in which there is only one available degree, in which the set of degrees available at the context is a singleton. Such contexts would appear to be ones in which a vague predicate was being used with a precise meaning, contrary to the ineluctability of vagueness. For example, if $\delta_{tall} = 6ft$, then the context $\sigma = \{\langle w, \delta_{tall}\rangle, \langle v, \delta_{tall}\rangle, \langle u, \delta_{tall}\rangle\}$ would be one in which the predicate "tall" was being used with the same meaning as the precise predicate "over exactly six feet in height." Furthermore, we may easily prove that at the context σ, the predicate "tall" has no borderline cases.

Proposition 2.1 There is no borderline case for "tall" at the context σ.

PROOF: Let α be an object. If $P(\alpha)$ is T or F at σ, then clearly α is not a borderline case for P at σ. So, suppose $P(\alpha)$ is I at σ. Now consider the context $\sigma' = \{\langle w, \delta_{tall}\rangle\}$. Clearly, the worlds available at σ' are a subset of those available at σ, the degrees available at σ' and σ are the same, and for any w in the worlds available at σ', if $\langle w, \delta_{tall}\rangle \in \sigma$ then $\langle w, \delta_{tall}\rangle \in \sigma'$. However, since σ' contains only one sequence, then for all sentences ϕ, ϕ is either T or F at σ'. Thus $P(\alpha)$ is not I at σ', and so α is not a borderline case for P.

The impossibility of contexts like σ looks to be evidence against our analysis.

Related to our analysis's blindness to the ineluctability of vagueness is a certain inadequacy of our Law of Borderlines. Theorem 2.6.4 states that for any standard-sensitive context dependent predicate P there exists a context σ in which there is a borderline case for P. This is, however, too weak. The ineluctability of vagueness requires that standard-sensitive predicates have borderline cases in *every* context. That we cannot with our current set-up derive this stronger law is witnessed by our theory's admittance of the context $\sigma = \{\langle w, \delta_{tall}\rangle, \langle v, \delta_{tall}\rangle, \langle u, \delta_{tall}\rangle\}$.

We must, then, give more serious thought to the nature of vagueness and try to understand why such precise contexts as σ are not possible. Ideally, this work will result in an extension of our formal analysis of vague predicates with which we may derive the following, stronger proposition as a theorem.

Ineluctability of Vagueness Let P be a standard sensitive context-dependent predicate, and O^P be its associated set of degrees. For any context $\sigma \subseteq W \times O^P$, there exists an object α such that α is a borderline case for P at σ.

3.2.2 The Strange Case of Higher Order Vagueness

The cascading anxieties and frustrations which encircle "higher order vagueness" usually mark the point at which semantic analyses of vague predicates grind to a halt, or at least begin to get very, very wobbly. The insecurity of one's footing is perhaps more the result of the problem's dizzying unclarity than any weakness of one's intellection. It is one of those puzzles in which the philosophical convictions are strong, and the seas of metaphor run high, but any linguistic data are nearly non-existent. One may legitimately wonder whether the "problem" isn't simply an artifact of abhorrently unrealistic gedankenexperiments. With all of this in mind, let us consider the matter.

The two-cent metaphor often used to state the problem of higher order vagueness is that "vague predicates don't have sharp boundaries *anywhere*." The problem can be a little more clearly stated as the fact that every vague predicate admits of a kind of infinite hierarchy of borderline cases. According to the standard description, not only can one imagine borderline cases for a vague predicate, but also *borderline* borderline cases. There are heights, for example, for which it is not clear whether it isn't clear whether it's tall. Once one accepts this "second order" vagueness, the existence of all higher orders fall like dominos. Borderline borderline borderline cases are imaginable, as well as borderline borderline borderline borderline cases and borderline borderline borderline ... and so on, *et cetera*, *ad nauseam*. The simplest way to state the problem, however, is that for any vague predicate Φ, the predicate "borderline case of Φ" is also vague.

There does not appear much concrete data in support of higher order vagueness, and one can't fight the feeling that the consensus that there is a problem is largely due to unconsciously interfering philosophical metaphors. Nonetheless, one cannot help but agree to intuitions like the following. English color words are vague, and there exist colors which are not properly described by any term in our color vocabulary. In some cases, we mitigate this weakness by inventing new color terms to grab hold of the borderline cases falling "in between" the terms we already have; thanks to Crayola, in fact, people make a living at this. Our inventing these new color words, however, can never eliminate that essential weakness of our color vocabulary, for the new color words always have their own borderline cases, and thinking up new words to cover *those* borderline cases seems not to help us in any fundamental way. We can never fill out our color vocabulary so completely that there would be no color word admitting of borderline cases. At the face of things, this seems to support the claim that the predicate "is a borderline case" must always itself admit of borderline cases.

The pressure is on, then, for an analysis of vagueness to explain what is going on behind these intuitions, and why they appear also to be an essential property of vague expressions. To this end, let us attempt to more precisely formulate the problem of higher order vagueness. Let I be our sentence operator meaning "it is indefinite whether", and let $I^n\phi$ be an abbreviation for the formula $\underbrace{I \ldots I}_{n}\phi$. One can state the problem of higher order vagueness as the fact that a dynamic semantics for vague predicates should be able to derive the following as a theorem.

Law of Higher Order Vagueness Let P be a vague predicate and O^P be its associated set of degrees. For any context $\sigma \subseteq W \times O^P$, and any natural number $n \geq 1$ there is some object α such that $I^nP(\alpha)$ is T at σ.

Before taking up this charge, however, it might be healthy to completely let off the steam which has already been occasionally moistening our discussion. Are there really an infinite number of "vagueness orders" for any vague predicate? I myself I have trouble taking seriously the concept of "borderline borderline borderline cases". That is, it is nigh impossible for me to imagine any, and even getting clear on what it would mean to be such a thing is difficult. This, of course, is due to the limits of my own imagination; perhaps the concept is quite comfortable to others. But, the clarity of anyone's intuitions runs out at some point, and after that point, who's to say how many orders of vagueness there are? Persons may have strong convictions on the matter, but that is likely due to their having already adopted some prejudices over what vagueness is. One's prior conceptions of the nature of vagueness fill in the details after the intuitions run thin. If you adopt the Fregean metaphor of vagueness being a "shadowy boundary", then the existence of the whole infinity of orders is natural. This metaphor, however, is as much a hypothesis about the semantics of vague predicates as the supervaluational analysis, and its predictions must not be confused with data.

Let us keep to the standard picture, however, and grant that any acceptable analysis of vague predicates must capture the Law of Higher Order Vagueness. The onus is now on our dynamic account to say something about these phenomena. Our dynamic account drops the ball; our system doesn't derive the law.

Proposition 2.2 Let P be a vague predicate and O^P be its associated set of degrees. There exists a context $\sigma \subseteq W \times O^P$ in which for any natural number $n \geq 1$ there is no object α such that $I^n P(\alpha)$ is T at σ.

PROOF: Let σ be the context $\{\langle w, \delta \rangle\}$, where $c(w) = \sigma$. Now, since σ consists of only one sequence, every sentence at σ will be either supported or rejected. So, let α be some object. We will prove by induction on n that for all $n \geq 1$, $I^n P(\alpha)$ is not T at σ.

For the base case, $P(\alpha)$ will be supported or rejected at σ, and so either $P(\alpha)$ is supported by σ or $\neg P(\alpha)$ is. Thus, $\langle w, \delta \rangle \models DP(\alpha)$ or $\langle w, \delta \rangle \models D\neg P(\alpha)$. Either way $\langle w, \delta \rangle \not\models \neg DP(\alpha) \wedge \neg D\neg P(\alpha)$, and so $\langle w, \delta \rangle \not\models IP(\alpha)$. Thus, $IP(\alpha)$ cannot be T at σ.

Suppose now that for all $m \leq n$, $I^m P(\alpha)$ is not T at σ; we will show that $II^n P(\alpha)$ is not T at σ. By assumption, $I^n P(\alpha)$ is not T at σ. Since σ contains only a single sequence, then $I^n P(\alpha)$ must be F at σ. By Proposition 2.9.17, if $I^n P(\alpha)$ is not I at σ, then $II^n P(\alpha)$ is not T at σ. Thus, $II^n P(\alpha)$ is not T at σ.

One reaction to this failing of our analysis, especially when combined with our manhandling of the ineluctability of vagueness, might be to reject our semantic theory of vague predicates as being at base unrealistic. A more conservative reaction, however, would locate the source of these troubles in our rather minimal formal theory of information change. In this thesis, we will, of course, adopt the more conservative stance. The greater part of this chapter will be devoted to an exploration of the origins of vagueness, pursuing the hunch that there may be a metaphysical hypothesis from which we could deduce some conditions holding on contexts in which vague predicates are used, conditions which might then permit our capture of the Ineluctability of Vagueness and the Law of Higher Order Vagueness. The hypothesis which we will converge upon portrays linguistic vagueness as being reducible to "vagueness" of a certain non-linguistic sort, and so the next section spells out in greater detail what such a hypothesis might look like.

We will end this section by noting two ways in which our work in Chapter 2 broke promises.

3.2.3 Broken Promises

We noted in Section 1.1 that standard-sensitive context dependency and the ability to have borderline cases are perfectly correlated; all predicates which are standard-sensitive admit of borderline cases and *vice versa*. We perhaps gave the impression at the beginning of Chapter 2 that there would be before page 55 some argument that our theory predicted this relationship, but none appeared. We were only able to show that all standard-sensitive predicates admit of borderline cases; we missed the "*vice versa*" part. Indeed, our analysis still at this point lacks the information required to show that predicates admitting of borderline cases, in the informal sense introduced in Section 1.1, must thereby be standard-sensitive. We would like, then, to further develop our analysis to account for the other direction of this relationship. We'll get around to it by the end of the chapter.

As a final reason for dissatisfaction with our present analysis, recall that it was said in Chapter 1 that the real objective of a theory of vagueness is to explain the properties of vague expressions by ascribing to them some underlying nature from which all their observed properties would follow. So far, though, we have not really attempted anything of the kind. Chapter 2 only proposed a dynamic semantics for standard-sensitive context dependent predicates with which we could predict their having borderline cases. We did not throw our hat in the ring and propose that vagueness *was* any one thing, nor did we point out a feature of vague predicates from which we could deduce their being standard-sensitive, having borderline cases and being susceptible to sorites paradoxes all in one. In this chapter, that will be corrected.

3.3 A Reductive Theory of Vagueness

In this section, I wish to make explicit a mish-mash of ideas and assumptions that underly the direction our discussion will soon take and which are important for locating the core claims of this thesis within a more general space of ideas. Everything in this section concerns how our work relates to that of others, nothing in this section seriously develops our overall argument, and so the reader is advised to skip ahead to Section 3.4 if he feels he could care less, really.

It would not be a distortion to say that metaphysics is that branch of philosophy which concerns questions such as which properties exist, which are real, and how different properties relate to one another. One subject of particular concern to metaphysicians is the relation of "reduction." What is meant, exactly, in saying that some property or fact "reduces" to another is a matter of fundamental and long standing debate. We might naively assume, however, that one property reduces to another if the nature or identity of that property can be explained in terms of the other property, but not *vice versa*.

No matter what methodology one adopts, if one seeks an answer to the question "what is x," then one is doing metaphysics. The question which drives our present study – "What is linguistic vagueness?" – is, then, a metaphysical question. We might, moreover, propose a "reduction" in order to answer this question. One such reduction would explain the existence in natural language of predicates exhibiting the combined properties of sorites-susceptibility, standard-sensitivity and having borderline cases as due to a property which does not depend for its existence on natural language having those predicates, but which also shares a great deal in common with them. Due to its superficial similarity to linguistic vagueness, this property may be called, for lack of a better homophone, "vagueness." A reduction of this sort we will call a "reductive theory of vagueness."

Reductive theories of vagueness are already witnessed within the philosophical literature. Graff 2000, for example, seeks to explain sorites-susceptibility by appeal to the "vagueness" of one's purposes. It argues that in some cases a person's *purposes* may be considered to be primitively "vague," in that they appear to be susceptible to sorites paradoxes, but their being so is independent of the precision of one's language. For example, I often have a purpose I might describe as "making some coffee." Certain states of affairs will clearly satisfy that purpose and others will clearly not. One drop of coffee, although it is "some coffee," will not satisfy my purposes; on the other hand, a full cup of coffee will. In between, however, there are states of affairs for which it isn't clear to me or anyone else whether my purposes would be satisfied [1]. Similarly, this purpose is subject to a sorites paradox: one drop of coffee clearly won't satisfy my purposes, and adding one drop of coffee to an amount which doesn't satisfy my purposes will not make an amount which does. In this sense, then, it seems right to characterize my purpose as "vague." However, as Graff 2000 persuasively argues, this "vagueness" of my purpose is not dependent on the vagueness of language. It could be hazarded that "make some coffee" is, strictly speaking, precise, but Graff 2000 does discuss more convincing examples.

Graff 2000 goes on to argue that, not only is this "vagueness" of purpose independent of linguistic vagueness, one might offer an analysis of vague predicates whereby their susceptibility to sorites paradoxes arises precisely because

[1] It should noted that Graff 2000 does not make this point, and explicitly wishes to avoid claiming that vague purposes may have borderline cases, though its reasons for doing so are not fully clear.

one's purposes can be "vague" in this newer sense. Thus, Graff 2000 would seek to explain, to reduce the vagueness of language to the "vagueness" of one's purposes. In the same vain, we will in this thesis extend our current semantic theory so that the vagueness of language is reduced to the "vagueness" of a particular entity, which we call a "vague selection." We will begin to grab hold of these things in the following section. It will be shown, moreover, that by doing so, we grab hold of a property of "vagueness" which is prior both to the vagueness of language and to the "vagueness" of one's purposes, as described by Graff.

3.4 Vague Selections

The perspective we shall defend may be summarized by the following observation. When I as an individual person use the vague predicate "tall," I have in mind some value, or range of values, marking a kind of lower bound on the heights of things I would describe as "tall." However, I would not say that this value or range of values is at all clear to me. To a large degree, I leave it open to myself how I will use the predicate "tall." For example, I may well get up in the morning and think "today, a person of roughly Sammy Sosa's height and above will be tall." Although the intention rarely finds itself in such an explicit form, all our uses of "tall" depend upon an intention somewhat akin to this exampling of Mr. Sosa. One could, perhaps, imagine me to be placed before the dense order of heights and forced to choose among them to determine my semantics for "tall." Faced with such a heavy decision, I flippantly wave my hands at the six-foot mark and say "uhm. . . there, roughly." Such a scene could very well be taken as a caricature of the hypothesis about to be offered.

Let us dwell, for the moment, on this supposed act of flippantly waving one's hands towards a position on some scale. I would like to claim that this act falls within a more general class of actions, one which has a widespread presence in human life as well as an independence from the ability for langauge to be vague. Consider, for the moment, an example borrowed from "Philosophical Investigations." As Wittgenstein points out, I can say something like "stand over there" while waving my hands before me, and everyone will fully understand what I mean. But, what do I mean when I say this? It seems that I make some decision about where one should stand; that much is obvious. However, it's false to say that for every area in space, I have decided whether your form should overlap it or not. Interestingly, it is *also* false to say that for every area of space, I have decided whether I have decided whether your form should overlap it or not. In fact, we find accompanying this act an infinite hierarchy of "uncertain" or "borderline" spaces quite similar to the infinite hierarchy of vagueness orders observed earlier for vague predicates.

Another example: you and I are talking about the fishing we did over the weekend. I mention that I caught a rather sizeable catfish, and you ask what, exactly, I mean by "sizeable." I hold out my hands some distance and say "'bout, yay big." You are satisfied, of course, but you could have pressed further. You might choose to be persnickety and ask, “By ‘yay big’, do you mean this big?” and indicate some length with your own hands. Now, depending on the length you indicate, I may reply back “sure, that big,” or “nah, little bigger.” There are other lengths, though, where it’s not fully clear to me whether I meant my original indication to cover them. These lengths, moreover, are not themselves precisely demarcated. There are lengths for which I may not be sure whether I’m unsure whether my original indication covered them, and so on until you get tired of counting.

Consider the act of asking someone to cut a pie or to cut a ribbon. You’ve just baked a pie, and you ask me from where I would like my slice. I’ll most likely reply back “lemme get that piece,” while making a little swooping motion across the pie surface with my index finger. In doing this, I do not decide for every crumb within the pie whether that crumb should be a piece of my slice, and nor do I decide for every crumb whether I decide for that crumb whether it’s in my slice. Instead, there are pie-portions which satisfy my indication, ones which do not, and ones for which I’m not certain. For that latter group, I may hesitate, perhaps, in accepting the slice. The uncertain slices, moreover, do not form a precise class. We could not exhaustively classify all pie portions as being “selected,” “not selected” or “uncertain,” and nor do we help matters by adding the categories “uncertainly uncertain” or “uncertainly uncertainly uncertain.”

The act of asking someone to cut a ribbon provides a greatly evocative mental image of the multi-layered uncertainty we are faced with in these scenarios. Suppose we are wrapping Christmas presents, and I ask you to cut a ribbon “there.” Most likely, your action will successfully satisfy my request, but this needn’t have been the case. There are points on the ribbon at which your cut would have displeased me. And, like in our other cases, there are points for which it is not clear to you or I whether the cut would have displeased me, for which it is not clear whether it isn’t clear whether the cut would have displeased me, and so forth. Where *sel* is a predicate holding of the points on our ribbon at which your cut would be acceptable, and I is our operator meaning “it is indefinite whether,” the following is a crude diagram of the epistemic state which follows my request.

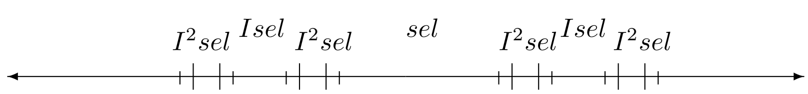

This picture is intended to represent that there are points, here marked by “*sel*,” at which your cut would satisfy my desire, as well as points (marked $Isel$) at which it is uncertain whether your cut would satisfy my desire, and points (marked I^2sel) at which it is uncertain whether it is uncertain whether your cut would satisfy my desire, and so on. Furthermore, this diagram captures the fact that these various classes of points have a particular *arrangement*. The “certainly selected” points are flanked on either side by intervals of “uncertainly selected" points. Each of these "uncertain" intervals is itself flanked by intervals of points for which it is not certain whether they are "uncertainly selected" points. Indeed, every "uncertainty order" is found to comprise a set of disjoint convex subsets of points, each such interval flanked on either side by intervals of one higher uncertainty order. The reader may notice here a striking similarity between the distribution of these "uncertainty orders" in our selection of points from the ribbon and the distribution of "vagueness orders" among the items arranged in a sorites series for a vague predicate [2]. The whole of this thesis may be condensed into the claim that this similarity is no coincidence, and it's the selection which has priority.

We have so far considered only acts in which a person is selecting or indicating a portion of physical space, but there are other acts of selection which are also saddled with these characteristic infinitely ascending levels of uncertainty. Consider the act of selecting a color from within the full continuous spectrum, or of choosing a pitch from within those frequencies audible to the human ear, or of indicating a particular moment in time. Indeed, selective acts of this sort show up in all walks of human life. Their ubiquity, moreover, lends the impression that their existence is independent of the ability for human language to be vague, simply for the reason that such acts as these don't necessarily depend upon the use of language. In fact, it could plausibly be claimed that *all* selective acts and decisions are infected with an infinitely unfolding uncertainty of this form [3], and so unless one wishes to hold that all acts of decision or selection depend upon the existence of language, one must conclude that the existence of these sorts acts does not depend upon language having any properties at all, let alone vagueness. As a final point in favor of the priority of these acts to linguistic vagueness, recall our example of the freshly baked pie. When I say to my pie-cutting friend "cut it here," the "vagueness," the indeterminacy of my statement intuitively depends not upon the vagueness of the verb "cut," but upon the vagueness of the word "here." I might have cut "cut" from my utterance without affecting the indeterminacy of my choice of pie-part. "Here," however, is not an inherently a vague word; "here" could be used to indicate a precisely demarcated area. In our imagined context, "here" is vague because my indication – my selection itself – was, in some sense, "vague."

Let us now dignify this collection of actions with a name.

Definition 4.1 (Vague Selection) A **vague selection** is an act through which some collection of objects forms a selection, and there exist objects for which it uncertain whether they are included in the selection, objects for which it is uncertain whether it is uncertain whether they are included in the selection, objects for which it is uncertain whether it is uncertain whether it is uncertain whether they are included in the selection, etc.

[2]Where P is the vague predicate, replace *sel* everywhere in the diagram above with "is a borderline case for P.

[3]Given the oft-made observation that nearly all langauge is vague, this provides a further parallel between the uncertainty of these selections and the vagueness of human language.

We contrast vague selections with precise selections.

Definition 4.2 (Precise Selection) A **precise selection** is an act through which some collection of objects forms a selection, and there exist no objects for which it is uncertain whether they are included in the selection, no objects for which it is uncertain whether it is uncertain whether they are included in the selection, no objects for which it is uncertain whether it is uncertain whether it is uncertain whether they are included in the selection, etc.

The notion of a "selection" used in the above two definitions is here left intentionally open, but its meaning should be more or less clear to the reader. Intuitively, the general act of making a "selection" encompasses such acts as reference and attending to something. Any act in which we might describe an object or class of objects as being "picked out" or "indicated" should be regarded as one in which a "selection" has been made. Let us also note here that the concept of a precise selection is not the complement of the concept of a vague selection. Consider a selection for which it is uncertain whether some points are included in the selection, but for every point it is certain whether it is uncertain whether it is included in the selection. By our definitions, such a selection would be neither a precise selection nor a vague selection. Therefore, if it were the case that all selections naturally occurring in human action were either vague or precise, it would be an interesting fact in need of explanation [4].

We argued above that the existence of vague selections is independent of the existence of linguistic vagueness. We might also argue that the existence of vague purposes, as described in Graff 2000, depends upon the existence of vague selections. All the examples of vague purposes put forth in Graff 2000 seem to depend, in some way, on a selection being vague. For example, the vagueness of my purpose in making a cup of coffee is due to the vagueness of my selection of coffee strength and amount. If I had made a precise selection of coffee size and strength, then my purpose in making some coffee would not be vague. Nevertheless, one might want to claim that every vague selection depends upon, or simply *is*, a vague purpose of some kind. This seems to me, however, to put quite a strain on the concept of a "purpose." Moreover, many vague selections appear to be instances of acts of reference, and the act of referring to something is arguably prior to the property of having a purpose. However, if it makes the reader more comfortable to understand a vague selection as issuing from a certain vague purpose, he may do so without it affecting his acceptance

[4] I would here like to indicate a potential weakness of Definition 4.1. There are certain acts which I believe we should want to classify as "vague selections" though they would not satisfy the condition put forth in 4.1. For example, given a line-up of, say, thirty men, I might say to someone standing next to me, "go talk to one of those guys," while waving my hand at some subsection from the line. Now, it seems that in such a scenario, not every one of the "uncertainty orders" of selected men is witnessed. Nevertheless, some of the orders *are* witnessed, and it intuitively feels that my selective act in this scenario is of the same kind as in the others considered above. This possible weakness of Definition 4.1, however, does not affect the material which follows.

of the primary claims of this section, nor will it affect the distinction between what we propose in this thesis and what is accomplished in Graff 2000.

Hopefully, the concept of a "vague selection" has by now been made fully clear. The thread which this whole chapter – perhaps this whole thesis – hangs from begins in the observation that one often *must* make a vague selection, even when attempting to be as precise as possible in one's intentions. One would perhaps like to think that a vague selection only occurs with rather loose intentions to refer, but the vagueness of one's selection can continue even when the most exacting intentions are mustered. Notice, for example, that one can never truly stand "exactly there," with full precision. There will always be variations in one's position which are simply irrelevant for the purposes of the agents involved. Similarly, close your eyes for the moment and confront yourself with the resulting homogenous void. Try, if you will, to indicate, to "direct your mind" towards some precise mathematical point in that great empty space. You will find that it is simply impossible. Likewise, when presented with a ruler, I may well try to indicate to others some point along its surface corresponding to one of the densely ordered values of customary length measurement. However, not only does my outer behavior seem to underdetermine to which point I am motioning, but in a real, non-skeptical sense so do my own intentions. That is, putting aside issues relating to the "radical indeterminacy of reference," one has even a certain first-person intuition that their own intentions simply fail to pick out a precise point along this ordering. I cannot even delude myself, it seems, into believing that there exists a particular point which is the object of my thoughts. Now, the reader may note that one can easily achieve the intended accuracy if his language is equipped with *names* for the elements of those dense orderings, a point which will be returned to shortly.

The fact that it is impossible to achieve such perfect precision in our thoughts was already known to the British Empiricists. However, the necessary indeterminacy of our thoughts in these domains takes on a familiar structure. Notice again that we strongly feel that, although our attention does not light upon a particular point in these orderings, it is certainly focused *somewhere*. Where, then, is our attention focused? It is fair to say that our intentions here make a vague selection just as they do when we wave our hands and intone "stand roughly there." When I close my eyes and focus on the resulting blackness, there are a number of points clearly falling within my range of attention. If one could somehow illuminate my eyelids to indicate points within that field, a considerable collection of them will be immediately recognized as ones to which I could or could not have referred. For other points, however, it may no longer be clear to me whether they fall within my referential intentions. Moreover, among these uncertain points there will be points on which I may not be sure whether I should truly regard them as points for which I am uncertain. Our imprecise intentions in these cases thus appear to be entities with precisely the kind of higher-order indeterminacy as the aforementioned vague selections. It is unclear, in fact, what important distinction might exist between these objects, and so it is proposed that the two simply be identified.

Although it may seem now that our use of vague selections is inescapable, the situation changes dramatically when our language is outfitted with names for the entities occupying these dense orderings. It is, for example, no great effort for me to focus my attention on a member of the dense ordering of rational numbers; all I need do is think of the number 3/4. Moreover, not every entity within the order need have a name. Thus, I can refer rather easily to a point amongst the real numbers simply by uttering the name π. Finally, the existence of numbers in the ordering is not essential either. There exists, for example, a dense linear ordering of all the sentences of the English language. We may, however, determinately refer to any member of that ordering simply by placing quotes around any sentence of our choosing.

All these considerations inspire the following principle, which we may take to be a psychological law.

Principle 4.3 (Selection Principle) The only way in which one may select objects from a dense ordering is through a vague selection or a precise selection. A precise selection, moreover, is only possible if there exist names in the language of the selector for some of the objects selected.

This Selection Principle is offered for now as simply a plausible hypothesis. No arguments will be put forth in support of it, but it should strike the reader as eminently true. I, for one, cannot imagine any way in which a person could make a selection of points from a dense order without that selection being either a vague or a precise one. Moreover, in all cases in which precise selections are imaginable, I have discovered that names for the points selected were necessary. One may rightly wonder why the Selection Principle is true, given that that it appears to be so robustly confirmed by one's own experience. We can speculate on this only to a certain extent, but it seems likely to be a very deep reflection both of our abilities to know our own judgments and of our inability to represent infinite sets that cannot be recursively "packed."

The main goals of this section were to informally introduce the idea of vague selections and to put in place the Selection Principle. The remainder of this chapter is an attempt to use vague selections and the Selection Principle to place constraints on our information states that will allow our derivation of the Law of Higher Order Vagueness. Our intended method of attacking higher order vagueness with vague selections, however, will require our semantic analysis to become quite serious about, and to depend quite strongly upon, the degree orderings for vague predicates introduced in the last chapter. Although this may strike one at first blush as intensely suspicious, in the next section I will say a few words in its favor.

3.5 Primitive Degrees

The account to be pursued in the remainder of this thesis will explain the vagueness of a predicate in terms of vague selections among the abstract degrees to which that predicate may hold. Thus, the "degree" of an object's ugliness is taken as prior to the vagueness of the predicate "ugly." However, there exists a strong intuition in some that this is quite suspect. Such persons would rather explain this mysterious concept of a "degree" in terms of a predicate's vagueness. I am myself not immune this intuition; the idea of leaving degrees prior to vagueness does infect me with the same kind of queasiness as I get when I begin taking seriously the concept of a "possible world." However, my queasiness is not an argument, and there don't seem to be many empirical reasons for not taking "coming-in-degrees" to be the more basic property. All that is really driving my distaste is the noble sentiment that we should unburden our theory of such a ponderous, medieval ontology. Moreover, upon further reflection, it seems that we don't truly purchase much ontological parsimony in holding vagueness as the more basic concept. The difficulty stems from what one is usually required to regard as "basic facts" in order to explain degrees away.

It appears that there is one particularly popular method for explaining the degrees to which a property holds in terms of a predicate's vagueness, one which has been regularly rediscovered by a few authors. Roughly, the idea is that since the extension of a vague predicate can vary without there being any change in the other properties of the objects in the domain, objects can be classified according to the number of contexts in which they appear within the extension of the predicate. In Kamp 1975, the "degree" an object has of a given property is the value of a measure function applied to the set of contexts in which the object falls within the extension of the predicate denoting the property. In the works of other philosophers, the domain is partitioned into equivalence classes based on the number of contexts in which the objects of the domain fall within the extension of the vague predicate. The "degree of P" which an object has is simply its P-equivalence class.

Both of these proposals take as basic and unexplained the ability of the vague predicate to "shift" its meaning from context to context. But, it's not immediately clear why this is more plausible as a "brute fact" than the ability for properties to come in degrees. Intuitively, the fact that I can on different occasions imbue "ugly" with different meanings relies on the fact that on different occasions what is an important *amount* of ugliness changes. Kamp 1975 even recognizes this intuition at several points in its exposition. When describing the contextual variation in the meaning of vague predicates (p. 139), it is suggested that this variation is the result of our "standards" for such predicates changing. In other sections, the paper mentions how our "criteria" for a vague predicate shift from one context to another. However, these informal notions of "standards" and "criteria" are nowhere explained, and nor do they enter into the formalism of the theory; for good reason, to be sure. If Kamp 1975 were to explain the notion of "degree" in terms of the contextual variation in a "criterion," it must then elucidate the distinction between a "degree" and a "criterion," which does not, for standard-sensitive predicates at least, appear to be either an appealing or a fruitful task.

Simply consider the adjective "big." My criteria for "big" include something like height and width and a proportional relationship between them. But is not the combination of these things merely something we would want to call a particular "degree of bigness?" That is, for any proposed criterion for the use of "big", it seems that we could equivalently phrase it in terms of something's degree of "bigness" being above a particular standard degree. Now, one in support of eliminating degrees might say, then, that we should here give up speaking in terms of "degrees" and instead speak primitively in terms of these criteria. But, this runs into the following snag: standard-sensitive predicates like "big" have criteria which *change* from context to context. This change, however, does not happen freely; one cannot simply begin using "big" in whatever way they like. It seems that the best way to capture this restriction on the variation of their use is to posit some predefined, structured space of criteria from which the speaker chooses at a particular context. However, once you have admitted to there being this kind of structured space of ready-made criteria, you may as well admit to speaking of degrees.

In summary, although one might want vagueness to be prior to "coming-in-degrees," standards should certainly be treated as prior to standard-sensitive context dependency, and there is no better way to model the "standards" and "criteria" required for our use of vague predicates than to model them as ordered sets of "degrees." Therefore, if there is a way to reduce coming-in-degrees to vagueness, it will not be by way of explaining degrees in terms of the standard-sensitive context dependency of vague predicates, such as in Kamp 1975 and other works.

Of course, a clever person may light upon some different method of explicating degrees in terms of vagueness, one which does not assume that the extension of vague predicates may shift recklessly between contexts [5]. However, since such a reduction cannot assume the standard-sensitivity of vague predicates, it may use only their ability to have borderline cases or to generate sorites paradoxes, and neither of these properties appears to be a useful tool for explaining what a "degree" is. But, of course, who knows? In the meantime, however, there are no clear benefits to taking a predicate's vagueness as prior to a property's coming in degrees, and toying with other possible lines of analysis leaves one with an uneasy worry that whatever form such an explication may take, the explicans will be quite a complex property, one which we would feel far better explicating in terms of the targeted explicandum.

In defense of our theory's appeal to a primitive realm of degrees, we may also here "break down the fourth wall," so to speak, and remind the reader that by our reasoning in Chapter 1, the most a theory of vagueness may accomplish is a sort of analogy or metaphor through which the various properties of vague

[5] Gaifman 2002, for example, attempts to identify an object α's degree of P with the higher-order vagueness predicates for P which may apply to α. This hypothesis, however, seems to be refuted by the following fact. Godzilla and King Kong are both "definitely large", "definitely definitely large", "definitely definitely definitely large", etc. That is, for all n, "definitely^n large" applies to both Godzilla and King Kong. Thus, any higher-order vagueness predicate for "large" will apply to Godzilla if and only if it applies to King Kong. However, it is possible for Godzilla to nevertheless be larger than King Kong. If "degrees of tallness" are to be employed in the semantics of the comparative in the standard way, then it would seem that Godzilla and King Kong can have differing degrees of tallness though all the same higher-order vagueness predicates apply to them.

predicates may be tied together and made sense of. Thus, one may understand our analysis in the following sections as simply one "picture" of the use of vague langauge, one that would demystify the co-occurrence of our three targeted properties. We aim not at "truth" in some broader sense, but at a metaphor which when put in mathematical terms provides greater results than competing metaphors. Although this sounds on paper like a sort of defeatist modesty, it is exactly the way in which many hypothesized entities enter into science and philosophy. The molecular theory of gasses, after all, was first offered by some of its developers as simply an analogy supporting some mathematical formulae which appeared to obtain results. Over time, such "metaphorical" objects tend to take on a greater reality, especially in the imaginations of non-specialists. That the same will happen for our hypothesized orderings of degrees is certainly doubtful, but the reader is here simply reminded that in the games of science and philosophy – and in particular linguistics – crazier things have happened.

3.6 The Reduction, Informally Described

What has anything we said in Section 3.4 to do with the semantics of vague language? In this section, we will begin uniting the Selection Principle and the concept of a vague selection with our dynamic analysis of vague predicates. The ideas laid out in this section will later receive a more formal treatment in the sections that follow. The result will be our ability to formally derive within our semantics both the Ineluctability of Vagueness and the Law of Higher Order Vagueness.

We will first make two linguistic observations. Their relevance is revealed in the third subsection.

3.6.1 Degree Orderings Are Dense

All vague predicates have the feature that they "come in degrees." Furthermore, it appears to be empirically well-supported that for any vague predicate, this ordering of degrees is *dense*. When speakers develop a conventional system of measurement for the properties denoted by these predicates, the values making up that system are always densely ordered. Think, for example, of inches, miles per hour, decibels and so forth. The density of the degree orderings also shows itself for vague predicates without associated systems of measurement. For all vague predicates, it is possible to order the domain of objects in virtue of how "well" they satisfy that predicate. Since we here take degrees as basic, we may characterize such orderings as being determined by the "degrees" of the predicate assigned to those objects. The density of the degree ordering may thus be inferred from the fact that within any such domain ordering we can, for any two differently ordered objects, imagine there being a third of intermediate location. Finally, it also appears that for any vague predicate, an object may vary continuously in the degree to which it satisfies that predicate[6].

[6]This is especially so for vague adjectives.

The implausibility of a discrete jump in the degree of a thing's "playfulness," for example, would support the notion that these degrees constitute a dense ordering.

Now, all these considerations may be struck down by finding a single, solid example of a vague expression whose degrees may be argued to form a well-ordering, but none have yet surfaced. Perhaps the clearest challenge to this claim comes from the vague adjective-quantifier "many." There is strong evidence to suggest that the degrees for "many" are cardinal numbers, or are at least isomorphic to the cardinal numbers, and so do not form a dense ordering. For one, the comparative form of "many" is "more." Our semantics for the comparative in Section 2.7, then, would seem to state that "there are more dogs than cats" is true if and only if the degree of "many" assigned to the set of dogs is greater than the degree of "many" assigned to the set of cats. However, as native English speakers, we also know that "there are more dogs than cats" is true just in case the cardinality of the set of dogs is greater than the cardinality of the set of cats. Another connection between cardinalities of sets and degrees of "many" is that one can, like any vague expression, order the domain of pluralities in virtue of how well they satisfy the adjective "many." Within this ordering, however, it seems that the domain of pluralities is divided into equivalence classes based on cardinality. Moreover, if one has a set of n objects and a set of $n + 1$ objects, it is impossible to find a plurality which can in such an ordering be fitted between them. If we hold that these orderings are determined by the degrees of "many" assigned to pluralities, this suggests at the very least that the ordering of "many-degrees" is isomorphic to the ordering of set cardinalities, and so is not a dense ordering.

On the face of things, then, "many" appears to be a counterexample to the claim that all vague expressions have a dense ordering of degrees. Other facts, however, indicate otherwise. There is evidence to suggest that "many" is underlyingly the same word as "much"; their only difference, it seems, lies in their phonological representation and in the fact that the former is syntactically restricted to count nouns, while the latter to mass nouns. Both words are adjective-quantifiers. Both words have the same comparative form, "more." Furthermore, there seems to be no real difference in their meaning, and in many languages both words are translated in the same way [7].

It is quite possible, then, that "many" and "much" are semantically the same expression, and differ only in their syntactic and phonological properties. If so, then the degrees of "many" would be the same as the degrees of "much." The degrees of "much," however, do appear to be densely ordered. Mass nouns are usually hypothesized to denote non-atomic groups, entities which when ordered according to their "size" form a dense ordering. Thus, if, as appears likely, the ordering of degrees of "much" is isomorphic to the ordering of sizes among non-atomic groups, we deduce that the degrees of "much" are densely ordered, and so, after all, the degrees of "many" are densely ordered as well.

[7] "Veel" in Dutch, "baie" in Afrikaans, and "viel" in German, for example, all translate as both "many" and "much."

The impression that the degrees of "many" are well-ordered may be explained away as a result of the syntactic restrictions governing "many." Count nouns are usually hypothesized to denote atomic groups, which for all intents and purposes we may as well consider sets. Suppose that degrees of "many/much" are assigned to sets in virtue of their cardinality. For every cardinality c there exists a degree δ_c of "many/much" such that all atomic groups of cardinality c are assigned δ_c as their degree of "many/much", and for any two cardinalities c, c', if $c < c'$ then $\delta_c <_{many/much} \delta_{c'}$. This would explain why when one orders a domain of pluralities denoted by count nouns, the ordering induced by their degrees of "many/much" is isomorphic to the ordering of their cardinalities. In general, it would predict that the subset of degrees of "many/much" assigned to the semantic values of count nouns forms, as a substructure of the full dense ordering of degrees of "many/much," a well-ordered set.

The following principle, then, stands up to falsification.

Principle 6.1 (Ordering Principle) Let P be a standard sensitive context dependent predicate and O^P be its set of degrees. The ordering $\langle O^P, <_P \rangle$ used in interpreting P is dense.

3.6.2 The Non-Synonymy Principle

Although its relevance here is perhaps obscure, let us take note of the following fact.

Principle 6.2 (Non-Synonymy Principle) If P is a standard-sensitive context dependent predicate and Q is not, then in all contexts, the meanings of P and Q will be distinct.

Given that a predicate which is not standard-sensitive is thereby precise, the Non-Synonymy Principle seems well supported; vague predicates never, in any context, take on the same meanings as other, precise predicates in the language. As to *why* the Non-Synonymy Principle should be true, one could answer in two ways. The first is that the Non-Synonymy Principle is simply a particular reflection of human language's general distaste for synonyms. It is often pointed out that in no human language does there exist any true "synonyms" [8], and that children in learning language seem likewise to assume that distinct words will have distinct meanings. Therefore, since predicates must always be distinct in their meaning, the standard-sensitive predicate "tall" can never take on the meaning of the distinct, non-standard-sensitive predicate "exactly six feet tall." A more likely explanation, however, is that the Non-Synonymy Principle arises from the speakers' pragmatic interests. Why would speakers ever in a context σ use the standard-sensitive word P to mean the same as the precise predicate Q?

[8] "Many" and "much" are no counterexamples to this claim if we suppose that their complementary distribution is in some way the result of a *sortal* rather than a syntactic restriction.

If they intend to express the meaning of Q, why not simply use Q? Endowing P with the meaning of Q in order to express that meaning appears to require a great deal of unmotivated extra work.

3.6.3 Standard Sensitivity and Vague Selections

Let us now finally reveal the relevance of vague selections and the Selection Principle to the semantics of standard-sensitive context dependent predicates.

First, we argue that an act of "making a selection" is essential to any use of a standard-sensitive predicate. Indeed, the grain of truth behind the implausible scenario in which I say to myself "today, everyone of roughly Sammy Sosa's height and above will be tall" is that the use of a standard-sensitive predicate requires one to select a standard for the predicate to which all objects will be compared in order to determine the predicate's applicability. At any occasion my use of the predicate "tall" requires my selecting some height or range of heights marking a lower bound on those objects I will describe as "tall." The selection of a standard is also necessary when negotiating with others over the meaning of a standard-sensitive predicate. If I were to enter into a conversation concerning "fat people," I might adjust my usage of the word "fat" to conform with that of the others in the conversation. However, contrary to the overly simple picture painted by our dynamic semantics, this process of adjustment in my usage seems to be fantastically complex, and consists of my making hypotheses concerning my company's standards for "fat" and testing those hypotheses against their future usage of the term. When my use of "fat" does not conform to the use of others in the conversation, I then *select* some other standard of fatness from the ordering of degrees, one that I hope more closely approximates the standard they have selected. The reader is invited to carry out an exhaustive consideration of the instances in which one uses a standard-sensitive context dependent predicate and to try to imagine a case in which the standards for the predicate have not, in a natural sense, been "selected" by the speakers. We believe that he won't be successful in this task [9].

This act of selecting a standard for a standard-sensitive predicate does not appear unlike those acts of selection we considered in Section 3.4. From the perspective of our dynamic semantics, the act of "selecting" a standard and "selecting" a color from within the full, continuous spectrum appear much the same. Recall that the set of standards with which one interprets a standard-sensitive predicate is simply the set of degrees associated with the predicate. Moreover, our Ordering Principle entails that any selection from this set of degrees will be a selection from a densely ordered set. We conclude then, by the Principle of Selection, that whenever a speaker uses a standard-sensitive predicate, he must make either a vague selection or a precise selection from the degrees associated with the predicate.

[9] An important set of facts which, again, our present study must leave unexamined concerns the restrictions on the selection which I may make for a predicate. The predicate "tall," for example, cannot have as its standard any height I like. These facts, which are related to the inferences discussed in footnote 19 of Chapter 2, are an important topic for future research.

However, the Principle of Selection also entails that the speaker can only make a precise selection of standards if his own language contains names for some of the selected standards. But, if the language contains names for those selected standards, then it will be possible for the speaker to form a complex predicate in his language using those names which has the same meaning within that context as his standard-sensitive predicate. For example, if in deciding upon a meaning for "tall" the speaker makes the precise selection of $\delta_{tall} = 6ft$, then his use of "tall" within the context will be exactly the same as "6ft in tallness or greater." Such uses as this are ruled out by the Principle of Non-Synonymy. Thus, the Principle of Non-Synonymy entails that a speaker, when faced with the task of selecting a standard for his use of the predicate, can only employ a *vague* selection. Recalling our motivation for that principle, we might say that speakers are prevented from using precise selections in determining the meanings of their standard-sensitive predicates because to do so would be a ridiculous extravagance. Why use the predicate "tall" with the meaning "6'2" and higher" when the speaker already possesses the means for expressing the latter concept much more efficiently?

Therefore, we deduce from the Selection Principle, the Ordering Principle, and the Non-Synonymy Principle that whenever one uses a standard-sensitive context dependent predicate, their use "depends" upon a vague selection of degrees from within the dense ordering associated with the predicate. We may state this condition upon contexts as the following.

Principle 6.3 (Vagueness Principle I) Let P be a standard-sensitive predicate, O^P be its associated densely ordered set of degrees, and $\sigma \subseteq W \times O^P$ be a context. There exists a vague selection of degrees from O^P determining the degrees available at σ.

The following section will be entirely taken up with spelling out more precisely what is meant by "determining" in the statement above. The intuitive meaning of "determining" here should, however, be clear to the reader. Principle 6.3 intuitively states that in choosing a standard with which to interpret a standard-sensitive predicate, the speaker makes a vague selection from among the degrees for the predicate. Thus, the degrees available in the context and the speaker's knowledge of what degrees are available in the context take on the kind of higher-order uncertainty witnessed in vague selections more generally.

We can already from this fact rather loosely reason our way to the Ineluctability of Vagueness and the Law of Higher Order Vagueness. In no vague selection of objects is a single object selected; always a vague selection from a dense ordering picks out some interval from within that ordering. Thus, our first Vagueness Principle would seem predict that there can be no context in which the set of available degrees is a singleton; no context exists in which the meaning of a standard-sensitive predicate is as sharp as the difference of a single degree of the predicate. Similarly, if the speaker's knowledge of what degrees lie within the set of those available at the context necessarily has the kind of higher-order uncertainty associated with vague selections, then it seems we might predict that in every use of a vague predicate, there are objects for which it is not certain in the context whether they are described by the predicate, ones for which it is not certain whether it isn't certain whether they are described by the predicate, and so on through all the infinite vagueness orders.

"Seems we might predict" and "we do predict," however, are two very different states of affairs. Our considerations above justify the intuitive, informal link between vague predicates and vague selections stated in our first Vagueness Principle: every use of a vague predicate depends upon the speaker making a vague selection amongst the degrees associated with the predicate. This is, of course, not yet enough for us to correctly derive from our semantic theory the Law of Higher Order Vagueness and the Ineluctability of Vagueness. In the next section, however, we develop a way of formally representing vague selections. With this system of representation we formulate a more precise statement of the first Vagueness Principle. Section 3.8 then reveals this statement to be sufficient for deriving our two targeted generalizations.

3.7 The Reduction, Formal

3.7.1 Formal Theory of Vague Selections

To mathematically model any phenomenon requires that one first isolate those properties of it which are important for one's study. Some mathematical object is then designed to represent any entity with those targeted properties. The intuition behind the present work is that for our study the most important feature of vague selections is the characteristic state of uncertainty which accompanies them. It is this nebulous state of uncertainty which above all else suggests the relationship between vague selections and vague predicates, and, as will soon be demonstrated, it is precisely this property of vague selections which is crucial for the truth of the Law of Higher Order Vagueness. Therefore, we will seek to represent vague selections as simply those peculiar epistemic states.

Such a representation, moreover, will happily be unambiguous. There appears to be a bijective relationship between vague selections and these epistemic states. That vague selections map onto these states may be seen from the fact that we consider such states as evidence for the presence of vague selections. That vague selections map into them follows from the plausible assumption that different vague selections result in different epistemic states of this sort.

We may, then, unambiguously identify a vague selection with an epistemic state from this class. So, to provide a formal representation of vague selections one need only provide a formal representation of these epistemic states. This last task may presumably be accomplished through the usual mathematical tools for representing epistemic states: pointed Kripke models. For a definition of these structures see, for example, Gerbrandy 1999 (p. 36) [10].

[10]To be precise, Gerbrandy 1999 defines pointed Kripke models for propositional modal logic, while we here use Kripke models for a first order modal logic. To save space, I will

In our representing vague selections as pointed Kripke models, however, we must engage in still further simplification. Our discussion in Section 3.4 reveals that the epistemic states characteristic of vague selections are themselves rather complex affairs. Consider the kind of vague selection of greatest interest to us: a vague selection from amongst a dense, linear ordering. Recall the "mental picture" of the resulting epistemic state, introduced in 3.4.

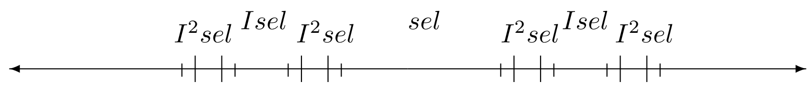

This diagram represents not only the existence of the infinitely ascending uncertainty orders, but also that these orders appear in a particular *arrangement*. Now, for the purposes of our present study, we will want to abstract from the particular arrangement of uncertainty levels depicted here. That is, it will be shown that in order to link vague selections to the structure of a context in such a way as to capture our targeted generalizations, it is sufficient for us to simply represent vague selections as epistemic states saddled with these unending strata of uncertainty orders. The particular arrangement of those orders will not be crucial for the establishing of our targeted result, and so we will ignore it in what follows.

Our goal, then, will be to simply define that class of pointed Kripke models which have all the infinite uncertainty orders of a vague selection. We may formulate our goal more precisely by introducing the predicate $sel(x)$, which ranges over the points in the dense linear ordering $\langle O^P, <_P \rangle$. The formula "$sel(\alpha)$" may be read as "the point α is included among the points selected from O^P." Let I, moreover, be our sentence operator read as "it is indefinite whether"; I will be provided a new model-based semantics shortly. Our mathematical models of vague selections will be pointed Kripke models taken from a class **vagsel** with the property that $\langle K, w \rangle \in$ **vagsel** if and only if for all $n \in \mathbb{N}$, there is some $\alpha \in O^P$ such that $\langle K, w \rangle \Vdash I^n sel(\alpha)$. This class, of course, will contain models which violate the characteristic arrangement of uncertainty orders depicted in our diagram. There will, for example, be a $\langle K, w \rangle \in$ **vagsel** and points $\alpha, \beta, \gamma, \delta, \epsilon \in O^P$ such that $\alpha <_P \beta <_P \gamma <_P \delta <_P \epsilon$ but $\langle K, w \rangle \Vdash Isel(\alpha)$, $\langle K, w \rangle \Vdash IIsel(\beta)$, $\langle K, w \rangle \Vdash Isel(\gamma)$, $\langle K, w \rangle \Vdash IIsel(\delta)$ and $\langle K, w \rangle \Vdash Isel(\epsilon)$, entailing, at the very least, that there are *three* 1-level uncertainty orders rather than two. Again, however, the characteristic distribution of the

make the reasonable assumption that the reader is quite familiar with such Kripke models and the standard definition of satisfaction in first order modal logic. To the more sophisticated readers: it should also be clear here that our formalism does not require commitments on such subtleties as whether the domains in these model structures should be constant, increasing, or what have you. Nevertheless, it may be convenient to stipulate that our models have constant domains, so that we may easily avoid the irrelevant complexities regarding modal identity, quantification across worlds, and the like.

uncertainty orders will be found not to be essential to our understanding of how their presence induces the Law of Higher Order Vagueness.

Defining the Class of Vague Selections

The goal stated above could, of course, be quite trivially met by just defining the class **vagsel** as all those pointed Kripke models $\langle K, w\rangle$ with the property that for all $n \in \mathbb{N}$, there is some $\alpha \in O^P$ such that $\langle K, w\rangle \Vdash I^n sel(\alpha)$. Such a move, however, would not nearly be as intellectually rewarding as providing a purely "structural" definition of **vagsel** and then proving that it has the desired semantic property. On the one hand, such a structural definition will provide some insight into how the members of **vagsel** appear, insight which will be important in reasoning about this class of structures in our proofs. On the other hand, it would, if done properly, allow us to prove that **vagsel** is a non-empty class of structures. Finally, we are prevented from offering such a definition here for the simple fact that we have not yet provided I with an explicit model-based semantics.

Let us then construct **vagsel** in the following way. Let $\langle O^P, <_P\rangle$ be our dense linear ordering of points. The class **vagsel** will always be relative to some such dense ordering. That is, for our purposes here, we are clearly interested only in vague selections over dense linear orderings of degrees, and so our formal representations of vague selections will reflect that preoccupation. One should, perhaps, think of the class **vagsel** as always implicitly indexed by some dense ordering $\langle O^P, <_P\rangle$, but here we will simply neglect this detail. We begin our construction with the definition of an admissible model.

Definition 7.1 (Admissible Model) A pointed Kripke model $\langle K, w\rangle$, where $K = (W, R, V)$ [11] is **admissible** if for all $u \in W$, there exists an open, convex non-trivial subset S of O^P such that for all $\alpha \in O^P$, $\langle K, u\rangle \Vdash sel(\alpha)$ iff $\alpha \in S$.

In the following, we always restrict ourselves to the class of admissible pointed Kripke models. This restriction plays an important role in our subsequent derivation of the Law of Higher Order Vagueness, and it also represents that one property of a vague selection is that the selector knows that he indicated *some* interval of points from amongst O^P.

Next, we define various classes of "uncertain" structures.

Definition 7.2 Where $\alpha \in O^P$, a pointed Kripke model $\langle K, w\rangle = \langle (W, R, V), w\rangle$ is **1-uncertain for** α if there are $v, v' \in W$ such that Rwv, Rwv', $\langle K, v\rangle \Vdash sel(\alpha)$ and $\langle K, v'\rangle \nVdash sel(\alpha)$. The class of pointed Kripke models 1-uncertain for α is denoted by α**-1-unc**. Where $\alpha \in O^P$, a pointed Kripke model $\langle K, w\rangle = \langle (W, R, V), w\rangle$ is **(n+1)-uncertain for** α if there are $v, v' \in W$ such that Rwv, Rwv', $\langle K, v\rangle \in \alpha$-n-unc, and $\langle K, v'\rangle \notin \alpha$-n-unc. The class of pointed Kripke

[11] By the coordinate "V", understand us here to always mean whatever formal machinery is required to associate each point in W with some interpretation of a first order language.

models (n+1)-uncertain for α is denoted by α-**(n+1)-unc**. A pointed Kripke model $\langle K, w\rangle$ is **n-uncertain** if there is some $\alpha \in O^P$ such that $\langle K, w\rangle \in \alpha$-n-unc. The class of n-uncertain pointed Kripke models is denoted by **n-unc**.

Although certainly a mouthful, Definition 7.2 is just what we need to formally represent a vague selection.

Definition 7.3 (Vagsel) The set of vague selections, **vagsel**, is the set containing all and only the admissible pointed Kripke models which are, for every $n \geq 1$, a member of n-unc; i.e., **vagsel** $= \bigcap\{$n-unc: $n \geq 1\}$.

That **vagsel** is non-empty can be seen by considering the following model. Let K be the model (W, R, V) where $W = \{w, v\}$ and $R = \{\langle w, w\rangle, \langle w, v\rangle\}$. Moreover, where S, S' are distinct, convex, open subsets of O^P, let $\langle K, w\rangle$ be such that for all $\alpha \in O^P$, $\langle K, w\rangle \Vdash sel(\alpha)$ iff $\alpha \in S$ and $\langle K, v\rangle$ be such that for all $\alpha \in O^P$, $\langle K, v\rangle \Vdash sel(\alpha)$ iff $\alpha \in S'$. The pointed model $\langle K, w\rangle$ can be pictured as below.

$$\begin{tikzcd} \circ w \arrow[loop left] \arrow[d] \\ \circ v \end{tikzcd}$$

Since $S \neq S'$, there is either some $\alpha \in S$ such that $\langle K, w\rangle \Vdash sel(\alpha)$ and $\langle K, v\rangle \nVdash sel(\alpha)$ or some $\alpha \in S'$ such that $\langle K, w\rangle \nVdash sel(\alpha)$ and $\langle K, v\rangle \Vdash sel(\alpha)$. Either way, since Rww and Rwv, there is some $\alpha \in O^P$ such that $\langle K, w\rangle \in \alpha$-1-unc. Now, since for no $w \in W$ does Rvw, then for no $n \in \mathbb{N}$ is $\langle K, v\rangle \in \alpha$-n-unc. Thus, quite clearly, $\langle K, w\rangle \in \alpha$-n-unc for all $n \geq 1$, and so is in n-unc for all $n \geq 1$. By definition, then, $\langle K, w\rangle \in$ **vagsel**.

With this definition of **vagsel**, we should now prove that $\langle K, w\rangle \in$ **vagsel** iff for all $n \in \mathbb{N}$, there is some $\alpha \in O^P$ such that $\langle K, w\rangle \Vdash I^n sel(\alpha)$. First, however, we must explicitly define our semantics for the operators I and D within a pointed Kripke model. Given their informal readings of "it definitely the case that" and "it is indefinite whether," the following presents itself as supremely natural.

Definition 7.4 Let $\langle K, w\rangle$ be a pointed Kripke model, where $K = (W, R, V)$. $\langle K, w\rangle \Vdash D\phi$ if for every $u \in W$ such that Rwu, $\langle K, u\rangle \Vdash \phi$. $\langle K, w\rangle \Vdash I\phi$ if there exists $u, v \in W$ such that Rwu and Rwv, $\langle K, u\rangle \Vdash \phi$, and $\langle K, v\rangle \nVdash \phi$.

The reader may see that Definition 7.4 preserves our earlier definition of $I\phi$ as $\neg D\phi \wedge \neg D\neg\phi$. Definition 7.3 and 7.4 are together enough to establish the following lemma [12].

[12] Both Lemma 7.5 and Theorem 7.6 assume that all models are admissible

Lemma 7.5 For all $\alpha \in O^P$ and $n \geq 1$, $\langle K, w\rangle \in \alpha$-n-unc iff $\langle K, w\rangle \Vdash I^n sel(\alpha)$.

PROOF: The proof is by induction on n. For the base case, consider that, by the semantics for I, $\langle K, w\rangle \Vdash Isel(\alpha)$ iff there are $v, u \in W$ such that Rwv, Rwu, $\langle K, v\rangle \Vdash sel(\alpha)$, and $\langle K, u\rangle \nVdash sel(\alpha)$, iff $\langle K, w\rangle \in \alpha$-1-unc. For the induction case, consider that $\langle K, w\rangle \Vdash I^{n+1}sel(\alpha)$, iff $\langle K, w\rangle \Vdash II^n sel(\alpha)$, iff, by the semantics of I, there are $v, u \in W$ such that Rwv, Rwu, $\langle K, v\rangle \Vdash I^n sel(\alpha)$ and $\langle K, u\rangle \nVdash I^n sel(\alpha)$, iff, by the Induction Hypothesis, $\langle K, v\rangle \in \alpha$-n-unc and $\langle K, u\rangle \notin \alpha$-n-unc, iff $\langle K, w\rangle \in \alpha$-(n+1)-unc.

Given this lemma, we can easily prove our desired theorem.

Theorem 7.6 (Correctness of Vagsel) $\langle K, w\rangle \in$ **vagsel** iff for all $n \in \mathbb{N}$, there is some $\alpha \in O^P$ such that $\langle K, w\rangle \Vdash I^n sel(\alpha)$.

PROOF: For the left-to-right direction, suppose that $\langle K, w\rangle \in$ **vagsel**, and that $n \in \mathbb{N}$. $I^0 sel(\alpha) = sel(\alpha)$, and so if $n = 0$, then the existence of an $\alpha \in O^P$ such that $\langle K, w\rangle \Vdash I^n sel(\alpha)$ follows from the admissibility of $\langle K, w\rangle$. Now suppose that $n \geq 1$. By definition of **vagsel**, there is some α such that $\langle K, w\rangle \in \alpha$-n-unc. By Lemma 7.5, $\langle K, w\rangle \Vdash I^n sel(\alpha)$. For the right-to-left direction, suppose that for all $n \in \mathbb{N}$, there is some $\alpha \in O^P$ such that $\langle K, w\rangle \Vdash I^n sel(\alpha)$. By Lemma 7.5, then, for all $n \geq 1$, there is some $\alpha \in O^P$ such that $\langle K, w\rangle \in \alpha$-n-unc. Thus, for all $n \geq 1$, $\langle K, w\rangle \in$ n-unc, and so $\langle K, w\rangle \in$ **vagsel**.

Theorem 7.6 shows that we have found a formal representation of vague selections which identifies them precisely as those epistemic states that have the unending strata of indeterminacy depicted in our diagram. In the next subsection, we relate these epistemic states to our information states.

3.7.2 Relationship Between Vague Selections and Contexts

We have developed a means for formally representing vague selections as pointed Kripke models. Now we must set up some way of relating these structures to the formal representations of contexts so essential to our dynamic semantics. In order to accomplish this in as easy and natural a way as possible, we are going to switch our semantic formalisms.

Up to this point, our dynamic semantics has employed what we called the "Stalnakerian Approach," representing contexts as sets of sequences. The work of Gerbrandy 1999 (Chapter 4) demonstrates that we may easily move from this way of representing information states to multi-dimensional pointed Kripke models. Let σ be a context for a conversation c occurring at world w in which a standard-sensitive predicate P measured by a dense ordering of degrees $\langle O^P, <_P\rangle$ is being used. Our present dynamic semantics represents σ as a subset of $W \times O^P$. A perfectly equivalent way of representing the knowledge of the speakers in c is by using a pointed multi-dimensional Kripke model $\langle K, \langle w, \delta \rangle \rangle$, where $K = (W', R', V')$, $W' \subseteq W \times O^P$, and for all $\langle w, \delta \rangle, \langle w', \delta' \rangle \in W'$, $R' \langle w, \delta \rangle \langle w', \delta' \rangle$ iff $\langle w', \delta' \rangle \in c(w)$. In this formalism, the knowledge of the speakers in the conversation c at world w is represented by those pairs $\langle w', \delta' \rangle$ accessible by R' from any pair $\langle w, \delta \rangle$ within W' containing w as a coordinate. This, then, allows us to represent the information states of our dynamic semantics with the structures standardly employed in epistemic logic. For example, with this new perspective, the context for the conversation c at world w, which we may have previously modelled as $\sigma = \{\langle w, \delta \rangle, \langle v, \delta \rangle\}$ where $c(w) = \sigma$ and $c(v) = \sigma$, could now be modelled as the pointed Kripke model $\langle K, \langle w, \delta \rangle \rangle$, where $K = (W', R', V')$, $W' = \{\langle w, \delta \rangle, \langle v, \delta \rangle\}$ and $R' = \{\langle \langle w, \delta \rangle, \langle w, \delta \rangle \rangle, \langle \langle w, \delta \rangle, \langle v, \delta \rangle \rangle, \langle \langle v, \delta \rangle, \langle v, \delta \rangle \rangle, \langle \langle v, \delta \rangle, \langle w, \delta \rangle \rangle\}$, pictured as below.

$$
\begin{tikzcd}
\langle w, \delta \rangle \arrow[loop below] \arrow[rrr, leftrightarrow] & & & \langle v, \delta \rangle \arrow[loop below]
\end{tikzcd}
$$

If these rather sketchy comments still leave the reader somewhat mystified as to our new, model-based representations of contexts, we suggest he take a quick look at Gerbrandy 1999, Chapter 4.

Throughout the rest of this thesis, we will use pointed Kripke models to represent contexts, as described in the manner above. This change in our formalism requires the redefinition of a few concepts which were introduced in the last chapter using the Stalnakerian Approach, and which will be important in the sections to follow. First, we would like our notion of "truth" within a context to remain largely unchanged.

Definition 7.7 Let P be a standard-sensitive context dependent predicate, let O^P be its densely ordered set of degrees, and let $\langle K, \langle w, \delta \rangle \rangle$, $K = (W', R', V')$, $W' \subseteq W \times O^P$, be a context. If ϕ is a sentence possibly containing P, then ϕ is **T** at $\langle K, \langle w, \delta \rangle \rangle$ if for all $\langle w', \delta' \rangle \in W'$ such that $R' \langle w, \delta \rangle \langle w', \delta' \rangle$, $\langle K, \langle w', \delta' \rangle \rangle \Vdash \phi$. ϕ is **F** at $\langle K, \langle w, \delta \rangle \rangle$ if for all $\langle w', \delta' \rangle \in W'$ such that $R' \langle w, \delta \rangle \langle w', \delta' \rangle$, $\langle K, \langle w', \delta' \rangle \rangle \nVdash \phi$. ϕ is **I** at $\langle K, \langle w, \delta \rangle \rangle$ otherwise.

Notice that by this definition, the "truth" of a sentence at a context is still a distinct notion from that sentence being modelled by a particular pair within the context. As concerns that latter relation, we again attempt to preserve as much from our earlier work as possible.

Definition 7.8 Let P be a standard-sensitive context dependent predicate, $\langle O^P, <_P \rangle$ be its dense ordering of degrees, f^P be its associated function, α be some object, and $\langle K, \langle w, \delta \rangle \rangle$, $K = (W', R', V')$, $W' \subseteq W \times O^P$, be a context.

$\langle K, \langle w, \delta \rangle \rangle \Vdash P(\alpha)$ iff $f^P(w)(\alpha) >_P \delta$.

We will, moreover, employ the model-based semantics for the operators D and I introduced in Definition 7.4. This now permits us to establish the following convenient fact.

Proposition 7.9 Let P be a standard-sensitive context dependent predicate, let O^P be its densely ordered set of degrees, let ϕ be a sentence possibly containing P, and let $\langle K, \langle w, \delta \rangle \rangle$, $K = (W', R', V')$, $W' \subseteq W \times O^P$, be a context. $\langle K, \langle w, \delta \rangle \rangle \Vdash I\phi$ iff ϕ is I at $\langle K, \langle w, \delta \rangle \rangle$.

PROOF: Follows trivially from Definition 7.4 and Definition 7.7.

Finally, we will have to redefine our notions of the set of degrees and the set of worlds available at a context. Again, we try to do this in as conservative a manner as possible.

Definition 7.10 Let P be a standard-sensitive predicate, O^P be its densely ordered set of degrees, and $\langle K, \langle w, \delta \rangle \rangle$, $K = (W', R', V')$, $W' \subseteq W \times O^P$ be a context. The **degrees available at** $\langle K, \langle w, \delta \rangle \rangle$ is the set $\{\delta' : \exists w' \in W \; R'\langle w, \delta \rangle \langle w', \delta' \rangle\}$. The **worlds available at** $\langle K, \langle w, \delta \rangle \rangle$ is the set $\{w' : \exists \delta' \in O^P \; R'\langle w, \delta \rangle \langle w', \delta' \rangle\}$.

We also adopt a new definition of "borderline case." Although the following is nearly identical in its prose form to our previous definition, it is understood to use our newer definitions of the notions of truth at a state and the degrees and worlds available at a state.

Definition 7.11 (Borderline Case) Let P be a standard-sensitive predicate, O^P be its associated ordered set of degrees, α be an object, and $\langle K, \langle w, \delta \rangle \rangle$, $K = (W', R', V')$, $W' \subseteq W \times O^P$ be a context. α is a **borderline case** for P at $\langle K, \langle w, \delta \rangle \rangle$ if $P(\alpha)$ is I at $\langle K, \langle w, \delta \rangle \rangle$, and $P(\alpha)$ is I at any information state $\langle K', \langle w', \delta' \rangle \rangle$, $K' = (W'', R'', V'')$, $W'' \subseteq W \times O^P$ such that the worlds available at $\langle K', \langle w', \delta' \rangle \rangle$ are a subset of the worlds available at $\langle K, \langle w, \delta \rangle \rangle$, the degrees available at $\langle K', \langle w', \delta' \rangle \rangle$ and $\langle K, \langle w, \delta \rangle \rangle$ are the same, and for any u in the set of worlds available at $\langle K', \langle w', \delta' \rangle \rangle$, if $R'\langle w, \delta \rangle \langle u, \gamma \rangle$, then $R''\langle w', \delta' \rangle \langle u, \gamma \rangle$.

We have now on the table a representation of information states as pointed Kripke models, as well as translation into this formalism of many of the key concepts introduced in the last chapter. Nevertheless, one of the main characters of our story has yet to appear. What of context change potentials and our interest in how vague language affects the context? In fact, our interest in the dynamics of vague language will largely fall by the wayside in the remainder of this thesis, and we will not spend any time at all in using this newer formalism to model the interpersonal use of vague language. However, Gerbrandy 1999

(Chapter 4) does show that we can represent the effect of updating such a context $\langle K, \langle w, \delta \rangle \rangle$, $K = (W', R', V')$ with the information that ϕ as the elimination from $\langle K, \langle w, \delta \rangle \rangle$ of any edges in R' connecting any pairs $\langle u, \gamma \rangle \in W'$ to pairs $\langle w', \delta' \rangle \in W'$ such that $\langle K, \langle w', \delta' \rangle \rangle \nVdash \phi$. Unfortunately, we will never make use of this theory of context change.

Now, the point of all this rigamarole was to allow us to relate vague selections as formally represented in the previous subsection with contexts in which standard-sensitive predicates are used. Following our account informally laid out in Section 3.6, this relationship amounts to the fact that in all contexts, there is some vague selection which determines the degrees available within the context. This informal statement, however, has yet to be given a formal interpretation. Let us now correct that with the introduction of the following Vagueness Principle.

Principle 7.12 (Vagueness Principle II) Let P be a vague predicate, O^P be its densely ordered set of degrees, and $\langle K, \langle w, \delta \rangle \rangle$ be a context such that $K = (W', R', V')$ where $W' \subseteq W \times O^P$. There is a vague selection $\langle K', w \rangle \in$ **vagsel** such that $K' = (W, R, V)$ and for all $\langle w, \delta \rangle \in W'$, $\{w' : \exists \delta' \in O^P \; R' \langle w, \delta \rangle \langle w', \delta' \rangle\} = \{w' : Rww'\}$ and $\{\delta' : \exists w' \in W \; R' \langle w, \delta \rangle \langle w', \delta' \rangle\} = \{\delta' : \langle K', w \rangle \Vdash sel(\delta')\}$.

The second Vagueness Principle satisfies one of our primary goals for this section: a law stating the relationship between vague selections as formally characterized and contexts as formally characterized in our dynamic semantics. Informally, it states that for any context in which a vague predicate is being used, there is a vague selection such that for any pair within that context, that vague selection determines which worlds within W and degrees within O^P are available in the context assigned to the conversation at the world making up the pair. Note that how these degrees happen to be paired with the worlds is left completely free by the Principle. In the next and final section of this chapter, we will prove that the Vagueness Principle above provides enough restriction on our set of contexts to prove the Law of Higher Order Vagueness and the Ineluctability of Vagueness. First, however, we have some bad news.

3.7.3 Downsizing Our Goals

As both the Law of Higher Order Vagueness and the Ineluctability of Vagueness are presently formulated, we are just asking too much from our dynamic semantics in demanding that it derive these generalizations. Simply put, given the way we've phrased them, both of our targeted generalizations are false. The Ineluctability of Vagueness, for example, literally states that in every context, every standard-sensitive predicate has borderline cases. But, that's just not true. As I look around me, there are currently *no* borderline cases for the standard-sensitive predicate "blue"; everything I can see is either clearly blue or clearly not blue. What we've fudged in our formulation of this generalization is that, intuitively, we want it to state that in every context there are *imaginable* borderline cases for any standard-sensitive predicate. The reason for our initial, sloppy formulation was that we hadn't yet offered any formal definition of what an "imaginable borderline case" is. Similarly, the Law of Higher Order Vagueness is currently stated too strongly. Take any standard-sensitive predicate you like. There certainly aren't, as I look around my office, objects to witness *every* vagueness order for that predicate. Again, what we would like our semantics to capture is the weaker claim that in every context, for every standard-sensitive predicate P and natural number $n \geq 1$, there is some *imaginable* object α such that $I^n P(\alpha)$.

How, though, can we represent there being in a context an imaginable borderline case for P or an imaginable witness for $I^n P$ without terribly complicating our theory of contexts? Let us start by extending our semantics of standard-sensitive predicates in such a way that they may also range over their own degrees.

Definition 7.13 Let P be a standard-sensitive predicate, $\langle O^P, <_P \rangle$ be its dense ordering of degrees, $\alpha, \delta \in O^P$, and $\langle K, \langle w, \delta \rangle \rangle$ be a context such that $K = (W', R', V')$ where $W' \subseteq W \times O^P$. $\langle K, \langle w, \delta \rangle \rangle \Vdash P(\alpha)$ if and only if $\alpha >_P \delta$.

Definition 7.13 could be viewed as an analysis of locutions such as "Six feet is tall" or "A BMI of 30 would definitely be fat." We would state, for example, that at a particular context "Six feet is tall" is true if and only if six feet happens to be taller than every degree available at the context. Of course, one may prefer a more sophisticated analysis of such locutions, perhaps one in which they are taken as "shorthand" for more complex propositions quantifying over the domain of individuals, but for our purposes here we choose to keep the analysis simpler.

Now, with Definition 7.13 in place, let us introduce the notion of a "borderline degree."

Definition 7.14 (Borderline Degree) Let P be a standard-sensitive predicate, O^P be its densely ordered set of degrees, $\alpha, \delta \in O^P$, and $\langle K, \langle w, \delta \rangle \rangle$ be a context such that $K = (W', R', V')$ where $W' \subseteq W \times O^P$. α is a **borderline degree** for P at $\langle K, \langle w, \delta \rangle \rangle$ if and only if $\langle K, \langle w, \delta \rangle \rangle \Vdash IP(\alpha)$.

We claim that the notion of a "borderline degree" captures effectively what is meant by an "imaginable borderline case" for a predicate. This point is really driven home by the following proposition.

Proposition 7.15 Let P be a standard-sensitive predicate, f^P be its associated function, O^P be its densely ordered set of degrees, $\alpha, \delta \in O^P$, β be some object, and $\langle K, \langle w, \delta \rangle \rangle$ be a context such that $K = (W', R', V')$ where $W' \subseteq W \times O^P$. If α is a borderline degree for P at $\langle K, \langle w, \delta \rangle \rangle$, and in all $\langle w', \delta' \rangle \in W'$ such that $R' \langle w, \delta \rangle \langle w', \delta' \rangle$, $f^P(w')(\beta) = \alpha$, then β is a borderline case for P at $\langle K, \langle w, \delta \rangle \rangle$.

PROOF: First, we will show that $P(\beta)$ is I at $\langle K, \langle w, \delta \rangle \rangle$. Since α is a borderline degree for P at $\langle K, \langle w, \delta \rangle \rangle$, by our semantics for I, there exist $\langle u, \gamma \rangle, \langle v, \epsilon \rangle \in W'$ such that $R' \langle w, \delta \rangle \langle u, \gamma \rangle$, $R' \langle w, \delta \rangle \langle v, \epsilon \rangle$, $\langle K, \langle u, \gamma \rangle \rangle \Vdash P(\alpha)$ and $\langle K, \langle v, \epsilon \rangle \rangle \nVdash P(\alpha)$. Thus, by Definition 7.13, $\alpha >_P \gamma$ and $\alpha \not>_P \epsilon$. Therefore, by assumption, $f^P(u)(\beta) >_P \gamma$ and $f^P(v)(\beta) \not>_P \epsilon$, and so $\langle K, \langle u, \gamma \rangle \rangle \Vdash P(\beta)$ and $\langle K, \langle v, \epsilon \rangle \rangle \nVdash P(\beta)$. By our semantics for I, then, $\langle K, \langle w, \delta \rangle \rangle \Vdash IP(\beta)$, and by Proposition 7.9, $P(\beta)$ is I at $\langle K, \langle w, \delta \rangle \rangle$.

Now, let $\langle K', \langle w', \delta' \rangle \rangle$, $K' = (W'', R'', V'')$, $W'' \subseteq W \times O^P$ be an information state such that the worlds available at $\langle K', \langle w', \delta' \rangle \rangle$ are a subset of the worlds available at $\langle K, \langle w, \delta \rangle \rangle$, the degrees available at $\langle K', \langle w', \delta' \rangle \rangle$ and $\langle K, \langle w, \delta \rangle \rangle$ are the same, and for any u in the set of worlds available at $\langle K', \langle w', \delta' \rangle \rangle$, if $R' \langle w, \delta \rangle \langle u, \gamma \rangle$, then $R'' \langle w', \delta' \rangle \langle u, \gamma \rangle$. We will show that $P(\beta)$ is I at $\langle K', \langle w', \delta' \rangle \rangle$. Now, since γ and ϵ are degrees available at $\langle K, \langle w, \delta \rangle \rangle$, they are also degrees available at $\langle K', \langle w', \delta' \rangle \rangle$. Thus, there are $\langle u', \gamma \rangle, \langle v', \epsilon \rangle \in W''$ such that $R'' \langle w', \delta' \rangle \langle u', \gamma \rangle$, $R'' \langle w', \delta' \rangle \langle v', \epsilon \rangle$. Moreover, since the worlds available at $\langle K', \langle w', \delta' \rangle \rangle$ are a subset of the worlds available at $\langle K, \langle w, \delta \rangle \rangle$, then we know that $f^P(u')(\beta) = \alpha$ and $f^P(v')(\beta) = \alpha$. Thus, $f^P(u')(\beta) >_P \gamma$ and $f^P(v')(\beta) \not>_P \epsilon$, and so $\langle K, \langle u', \gamma \rangle \rangle \Vdash P(\beta)$ while $\langle K, \langle v', \epsilon \rangle \rangle \nVdash P(\beta)$. Therefore, $\langle K', \langle w', \delta' \rangle \rangle \Vdash IP(\beta)$, and so again by Proposition 7.9, $P(\beta)$ is I at $\langle K', \langle w', \delta' \rangle \rangle$.

Proposition 7.15 establishes that if α is a borderline degree for P in some context, then any object for which it's known that its degree of P is α will be a borderline case for P at that context. Now, what precisely does one do when one imagines a possible borderline case for a predicate P? Intuitively, one pictures in their mind some object whose degree of P is precisely known to them, but yet there is an indeterminacy in whether P applies to it. One way of characterizing this act is that for some borderline degree α for P, one pictures an object for which it's known that its degree of P is α. That is, in my imagining a borderline case for P, I find the "gap" for P and imagine to myself an object whose degree of P falls precisely in that gap. Clearly, if such an object were to exist, then by Proposition 7.15 there would be a borderline case for P. Thus, it seems that for there to be an imaginable borderline case for P at some context it is sufficient that there be at that context a borderline degree for P. Furthermore, it's clear that the existence of an imaginable borderline case for P implies the presence of a borderline degree for P; if in some context there were no borderline degrees for P, then P must in that context be used with a precise sense, and so intuitively there could be no imaginable borderline cases for P.

Thus, for the notion of an "imaginable borderline case for P," we shall throughout this thesis substitute that of a "borderline degree for P". Because of our equation of imaginable borderline cases and borderline degrees, we can now formulate the following as our desired results.

Law of Higher Order Vagueness Let P be a vague predicate and O^P be its densely ordered set of degrees. For any context $\langle K, \langle w, \delta\rangle\rangle$ in which P is being used, and any natural number $n \geq 1$ there is some degree $\alpha \in O^P$ such that $\langle K, \langle w, \delta\rangle\rangle \Vdash I^n P(\alpha)$.

Ineluctability of Vagueness Let P be a vague predicate and O^P be its densely ordered set of degrees. For any context $\langle K, \langle w, \delta\rangle\rangle$ in which P is being used, there is some degree $\alpha \in O^P$ such that $\langle K, \langle w, \delta\rangle\rangle \Vdash IP(\alpha)$.

Clearly, when formulated in this way, the Law of Higher Order Vagueness implies the Ineluctability of Vagueness. Therefore, in order to overcome the challenges raised in Sections 3.2.1 and 3.2.2, it is sufficient that we prove the Law of Higher Order Vagueness. In the next section, we do just that.

3.8 Overcoming Our Adversities

3.8.1 Proving the Law of Higher Order Vagueness

In order to prove the Law of Higher Order Vagueness, we must first establish a number of lemmas.

Lemma 8.1 Let P be a vague predicate, O^P be its densely ordered set of degrees, $\alpha \in O^P$, $\langle K, \langle w, \delta\rangle\rangle$ be a context such that $K = (W', R', V')$ where $W' \subseteq W \times O^P$, and $\langle K', w\rangle$ be an admissible pointed Kripke model such that $\{\delta' : \exists w' \in W\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\} = \{\delta' : \langle K', w\rangle \Vdash sel(\delta')\}$. $\langle K, \langle w, \delta\rangle\rangle \Vdash IP(\alpha)$ iff $\alpha \in \{\delta' : \exists w' \in W\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\}$.

PROOF: $\Leftarrow$ Suppose that $\alpha \in \{\delta' : \exists w' \in W\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\}$. By assumption, $\langle K', w\rangle$ is an admissible model, and so $\{\delta' : \langle K', w\rangle \Vdash sel(\delta')\}$ is a convex, open, non-trivial subset of O^P. Thus, $\{\delta' : \exists w' \in W\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\}$ is a convex, open, non-trivial subset. Therefore, there is a $\alpha' \in O^P$ such that $\alpha' <_P \alpha$ and $\alpha' \in \{\delta' : \exists w' \in W\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\}$. So suppose that $R'\langle w, \delta\rangle\langle w', \alpha'\rangle$ and $R'\langle w, \delta\rangle\langle w'', \alpha\rangle$. Since $\alpha' <_P \alpha$, $\langle K, \langle w', \alpha'\rangle\rangle \Vdash P(\alpha)$. However, since $\alpha \not<_P \alpha$, $\langle K, \langle w'', \alpha\rangle\rangle \nVdash P(\alpha)$. Thus, by our semantics for I, $\langle K, \langle w, \delta\rangle\rangle \Vdash IP(\alpha)$.

$\Rightarrow$ Suppose that $\alpha \notin \{\delta' : \exists w' \in W\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\}$. Now, since $\{\delta' : \exists w' \in W\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\}$ is a convex, non-trivial, open subset, α must be either a proper upper bound or a proper lower bound for that set. Suppose α is a proper lower bound for the set, and $\langle w', \delta'\rangle$ is such that $R'\langle w, \delta\rangle\langle w', \delta'\rangle$. Since $\alpha <_P \delta'$, then $\langle K, \langle w', \delta'\rangle\rangle \nVdash P(\alpha)$. Since $\langle w', \delta'\rangle$ was arbitrary, $\langle K, \langle w, \delta\rangle\rangle \nVdash IP(a)$. Similarly, suppose α is a proper upper bound for the set and $\langle w', \delta'\rangle$ is such that $R'\langle w, \delta\rangle\langle w', \delta'\rangle$. Since $\alpha >_P \delta'$, then $\langle K, \langle w', \delta'\rangle\rangle \Vdash P(\alpha)$. Thus, again, $\langle K, \langle w, \delta\rangle\rangle \nVdash IP(a)$. Therefore, we conclude that $\langle K, \langle w, \delta\rangle\rangle \nVdash IP(a)$.

Lemma 8.2 Let P be a vague predicate, O^P be its densely ordered set of degrees, $\alpha \in O^P$, $\langle K, \langle w, \delta\rangle\rangle$ be a context such that $K = (W', R', V')$ where

$W' \subseteq W \times O^P$, and $\langle K', w\rangle$ be an admissible pointed Kripke model such that for all $\langle w, \delta\rangle \in W'$, $\{w' : \exists\delta' \in O^P\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\} = \{w' : Rww'\}$ and $\{\delta' : \exists w' \in W\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\} = \{\delta' : \langle K', w\rangle \Vdash \mathit{sel}(\delta')\}$. For all $n \geq 1$, $\langle K, \langle w, \delta\rangle\rangle \Vdash I^{n+1}P(\alpha)$ iff $\langle K', w\rangle \in \alpha$-n-unc.

PROOF: The proof will be by induction on n.

Proof of the claim for n = 1

$\Leftarrow$ Suppose that $\langle K', w\rangle$ is 1-uncertain for α. Then, by definition, there are $v, u \in W$ such that Rwv, Rwu, $\langle K', v\rangle \Vdash \mathit{sel}(\alpha)$ and $\langle K', u\rangle \nVdash \mathit{sel}(\alpha)$. Thus, $\alpha \in \{\delta' : \langle K', v\rangle \Vdash \mathit{sel}(\delta')\}$ and $\alpha \notin \{\delta' : \langle K', u\rangle \Vdash \mathit{sel}(\delta')\}$. Now, since $\{w' : \exists\delta' \in O^P\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\} = \{w' : Rww'\}$ there are $\gamma, \epsilon \in O^P$ such that $R'\langle w, \delta\rangle\langle v, \gamma\rangle$ and $R'\langle w, \delta\rangle\langle u, \epsilon\rangle$. Furthermore, since $\{\delta' : \exists w' \in W\ R'\langle v, \gamma\rangle\langle w', \delta'\rangle\} = \{\delta' : \langle K', v\rangle \Vdash \mathit{sel}(\delta')\}$ and $\{\delta' : \exists w' \in W\ R'\langle u, \epsilon\rangle\langle w', \delta'\rangle\} = \{\delta' : \langle K', u\rangle \Vdash \mathit{sel}(\delta')\}$, $\alpha \in \{\delta' : \exists w' \in W\ R'\langle v, \gamma\rangle\langle w', \delta'\rangle\}$ and $\alpha \notin \{\delta' : \exists w' \in W\ R'\langle u, \epsilon\rangle\langle w', \delta'\rangle\}$. Now, since $\langle K', w\rangle$ is admissible, then $\langle K', v\rangle$ and $\langle K', u\rangle$ are admissible. Thus, if $\alpha \notin \{\delta' : \exists w' \in W\ R'\langle u, \epsilon\rangle\langle w', \delta'\rangle\}$, then by Lemma 8.1, $\langle K, \langle u, \epsilon\rangle\rangle \nVdash IP(\alpha)$. Likewise, if $\alpha \in \{\delta' : \exists w' \in W\ R'\langle v, \gamma\rangle\langle w', \delta'\rangle\}$, then, $\langle K, \langle v, \gamma\rangle\rangle \Vdash IP(\alpha)$. Since, by assumption, $R'\langle w, \delta\rangle\langle v, \gamma\rangle$ and $R'\langle w, \delta\rangle\langle u, \epsilon\rangle$, given our semantics for I, we conclude that $\langle K, \langle w, \delta\rangle\rangle \Vdash IIP(\alpha)$.

$\Rightarrow$ Suppose that $\langle K, \langle w, \delta\rangle\rangle \Vdash IIP(\alpha)$. By our semantics for I, this entails that there are $\langle v, \gamma\rangle, \langle u, \epsilon\rangle \in W'$ such that $R'\langle w, \delta\rangle\langle v, \gamma\rangle$, $R'\langle w, \delta\rangle\langle u, \epsilon\rangle$, $\langle K, \langle v, \gamma\rangle\rangle \Vdash IP(\alpha)$ and $\langle K, \langle u, \epsilon\rangle\rangle \nVdash IP(\alpha)$. Again, since $\langle K', w\rangle$ is admissible, then $\langle K', v\rangle$ and $\langle K', u\rangle$ are admissible. Thus, by Lemma 8.1, we may conclude that $\alpha \in \{\delta' : \exists w' \in W\ R'\langle v, \gamma\rangle\langle w', \delta'\rangle\}$ and $\alpha \notin \{\delta' : \exists w' \in W\ R'\langle u, \epsilon\rangle\langle w', \delta'\rangle\}$. Therefore, $\alpha \in \{\delta' : \langle K', v\rangle \Vdash \mathit{sel}(\delta')\}$ and $\alpha \notin \{\delta' : \langle K', u\rangle \Vdash \mathit{sel}(\delta')\}$. Thus, $\langle K', v\rangle \Vdash \mathit{sel}(\alpha)$ and $\langle K', u\rangle \nVdash \mathit{sel}(\alpha)$. From the fact that $\{w' : \exists\delta' \in O^P\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\} = \{w' : Rww'\}$, we conclude that Rwv and Rwu. By definition, then, $\langle K', w\rangle$ is 1-uncertain for α.

Proof of the claim for n = m+1

$\langle K', w\rangle$ is n-unc for α iff there are $v, u \in W$ such that Rwv, Rwu, $\langle K', v\rangle \in \alpha$-m-unc and $\langle K', u\rangle \notin \alpha$-m-unc, iff, by our induction hypothesis and the fact that $\{w' : \exists\delta' \in O^P\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\} = \{w' : Rww'\}$, for some $\gamma, \epsilon \in O^P$, $\langle K, \langle v, \gamma\rangle\rangle \Vdash I^nP(\alpha)$ and $\langle K, \langle u, \epsilon\rangle\rangle \nVdash I^nP(\alpha)$, iff, by our semantics for I, $\langle K, \langle w, \delta\rangle\rangle \Vdash II^nP(\alpha)$ iff $\langle K\langle w, \delta\rangle\rangle \Vdash I^{n+1}P(\alpha)$.

Lemma 8.3 Let P be a vague predicate, O^P be its densely ordered set of degrees, $\langle K, \langle w, \delta\rangle\rangle$ be a context such that $K = (W', R', V')$ where $W' \subseteq W \times O^P$, and $\langle K', w\rangle$ be an admissible pointed Kripke model such that $\{\delta' : \exists w' \in W\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\} = \{\delta' : \langle K', w\rangle \Vdash \mathit{sel}(\delta')\}$. There exists some $\alpha \in O^P$ such that $\langle K, \langle w, \delta\rangle\rangle \Vdash IP(\alpha)$.

PROOF: Since $\langle K', w\rangle$ is admissible, we know that $\{\delta' : \exists w' \in W\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\}$ is a convex, open, nontrivial subset of O^P. By its non-triviality, there is some

$\alpha \in O^P$ such that $\alpha \in \{\delta' : \exists w' \in W\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\}$. Thus, by Lemma 8.1, $\langle K, \langle w, \delta\rangle\rangle \Vdash IP(\alpha)$.

Now we can finally derive the Law of Higher Order Vagueness.

Theorem 8.4 (Law of Higher Order Vagueness) Let P be a vague predicate, O^P be its densely ordered set of degrees, and $\langle K, \langle w, \delta\rangle\rangle$ be a context such that $K = (W', R', V')$ where $W' \subseteq W \times O^P$. For all $n \geq 1$, there is some degree $\alpha \in O^P$ such that $\langle K, \langle w, \delta\rangle\rangle \Vdash I^n P(\alpha)$.

PROOF: By Principle 7.12 (Vagueness Principle II), there is a vague selection $\langle K', w\rangle \in$ **vagsel** such that $K' = (W, R, V)$ and for all $\langle w, \delta\rangle \in W'$, $\{w' : \exists \delta' \in O^P\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\} = \{w' : Rww'\}$ and $\{\delta' : \exists w' \in W\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\} = \{\delta' : \langle K, w\rangle \Vdash sel(\delta')\}$. Now suppose that $n \geq 1$. Since $\langle K', w\rangle \in$ **vagsel**, there is some $\alpha \in O^P$ such that $\langle K', w\rangle \in \alpha$-n-unc. Moreover, since $\langle K', w\rangle \in$ **vagsel**, $\langle K', w\rangle$ is admissible. Therefore, by Lemma 8.2, $\langle K, \langle w, \delta\rangle\rangle \Vdash I^{n+1}P(\alpha)$. Thus, for all $n \geq 1$, there exists an $\alpha \in O^P$ such that $\langle K, \langle w, \delta\rangle\rangle \Vdash I^{n+1}P(\alpha)$. Finally, from this and Lemma 8.3 follows the full Law of Higher Order Vagueness.

From this we now obtain the Ineluctability of Vagueness.

Corollary 8.5 (Ineluctability of Vagueness) Let P be a vague predicate, O^P be its densely ordered set of degrees, and $\langle K, \langle w, \delta\rangle\rangle$ be a context such that $K = (W', R', V')$ where $W' \subseteq W \times O^P$. There is some degree $\alpha \in O^P$ such that $\langle K, \langle w, \delta\rangle\rangle \Vdash IP(\alpha)$.

PROOF: Trivial consequence of Theorem 8.4.

We have thus rescued our dynamic semantics of standard-sensitive predicates from the challenges raised in Section 3.2. However, we have not yet lived up to the promises broken in the previous chapter. In the following subsections, we finally make good.

3.8.2 Borderlines Equivalent to Standards

Corollary 8.5 firmly establishes that our theory of dynamic semantics predicts all standard-sensitive predicates to have imaginable borderline cases in all contexts. We would like now for our theory to predict that any predicate for which there is at least *one* context in which it has a borderline case is thereby standard-sensitive context dependent.

However, establishing this direction of the relationship between standard-sensitivity and the ability to have borderline cases is for us at the moment not possible. From the fact that a predicate P is not standard-sensitive, we wish to deduce that it cannot have borderline cases *in any context*. Now, if we were to simply stipulate a “borderline case” to be an object satisfying our definition in 7.11, then this result would follow quite easily. But, we are after richer and more interesting game. We would like to rule out there being for P anything satisfying our informal notion of a "borderline case" introduced in Section 1.1: an entity for which one can't say whether P holds of it, and no matter what one learns about the world could change that. In the pursuit of this goal, however, we become stuck in whatever train of reasoning we take, because we have not yet explicitly defined our language. Without an explicit definition of our language, we simply do not have enough on the table in order to obtain this result. That is, without our making some explicit, constraining assumptions regarding the possible forms of predication which can occur in human language, we cannot rule out the existence of bizarre, undreamt of ways for there to be in some context an object for which it isn't known whether the predicate holds of it and this ignorance is not the result of one's lack of factual knowledge. Simply put, from the fact that P is not standard-sensitive, we cannot yet conclude anything about its potential to have borderline cases, because we can conclude nothing about its semantics.

Therefore, we are lead to the conclusion that in order to meet our goal of predicting the full relationship between standard-sensitivity and borderline cases, we must to a certain extent stipulate there being a dichotomy in our language between vague and precise predicates. The way we shall do this is by setting up within our semantic system only two ways in which one may use a predicate: all predicates in our language are either standard-sensitive or they are assigned as their semantic value classical intensions. Let us state that for any language $\mathcal{L}$, there are two classes of predicates **vague** and **precise**, that every predicate P of $\mathcal{L}$ is a member of exactly one of these classes, and that the semantics of P is entirely determined by which of these classes it is a member of. This last assumption is further elaborated through the introduction of the following two semantic principles.

Principle 8.6 (Principles of Predication)

1. $\forall P \in$ **vague**, $\langle K, \langle w, \delta \rangle \rangle \Vdash P\alpha$ iff $f^P(w)(\alpha) >_P \delta$

2. $\forall P \in$ **precise**, $\langle K, \langle w, \delta \rangle \rangle \Vdash P\alpha$ iff $\alpha \in \mathbf{P}(w)$.

These principles state that all vague predicates are merely standard-sensitive context dependent ones, and that all non-vague predicates are interpreted as intensions, functions from possible worlds to sets of objects. Now we wish to show that for all non-vague predicates there is no context in which they have a borderline case.

Proposition 8.7 (Vagueness is Standard-Sensitivity I) Let $P \in$ **precise**, $\langle K, \langle w, \delta \rangle \rangle$ be any context, and α be any object. α is not a borderline case for P in $\langle K, \langle w, \delta \rangle \rangle$.

ARGUMENT: In returning to our informal notion of a borderline case, we've stepped somewhat outside the realm of precise proof, but the following argument should satisfy the reader. Suppose that α is a borderline case for P at our context $\langle K, \langle w, \delta \rangle \rangle$. Our informal notion of a borderline case at $\langle K, \langle w, \delta \rangle \rangle$ states that no matter what one learns about the world, one should be unable to say whether $P(\alpha)$. Thus, even if one had the knowledge of a god, one couldn't say whether $P(\alpha)$. Formally, we may think of the knowledge of a god as being a state just like $\langle K, \langle w, \delta \rangle \rangle$, but in which there is only one available world, v. Now, by our second Vagueness Principle such a context is, strictly speaking, not possible; however, we're pretending to be gods here, not people. So, if one still cannot say at this divine context whether $P(\alpha)$, it must be because there are other possibilities still open at this state regarding the truth of $P(\alpha)$. But, all of these possibilities contain the same possible world, v. Thus, if at least two of these possibilities differ over the verity of $P(\alpha)$, it cannot be that P is a member of **precise**. If P were a member of **precise**, then, by Principle 8.6, P's extension would the same for all the remaining possibilities, and so those possibilities couldn't differ in their acceptance of $P(\alpha)$. So, $P \notin$ **precise**.

We claim, then, that through our postulation of the dichotomy between the classes of **vague** and **precise**, we capture the fact that a predicate admits of borderline cases if and only if it is standard-sensitive.

3.8.3 On the Nature of Vague Predicates

In the previous subsection, we finally made good on our promise to promote a semantic analysis explaining the necessary relationship between standard-sensitivity and the possibility of having borderline cases. We also inadvertently made good on our promise to offer some hypothesis regarding the "nature of vagueness." Proposition 8.7 permits us to venture that vagueness simply *is* standard-sensitive context dependency.

Now, the claim that vagueness is standard-sensitivity, or at least that vagueness and standard-sensitivity have a deliciously sordid relationship, is one which has been made in the literature before [13]. In order for this accusation to stick, however, it must be argued that standard-sensitive context dependency is both necessary and sufficient for the other two properties of vague predicates: the ability to have borderline cases and the susceptibility to sorites paradoxes. To my knowledge, no such arguments yet appear in the literature. In this chapter, however, we were able to argue, thanks to the added assistance of our theory of vague selections and our dichotomy between the classes **vague** and **precise**, that standard-sensitivity is necessary and sufficient for a predicate's ability to have borderline cases. In the next chapter, we will argue that, moreover, our theory of vague selections is enough to guarantee us that standard-sensitive context dependency is necessary and sufficient for a predicate's being susceptible to sorites paradoxes. This will complete our argument that the underlying na-

[13] See, for example, Bosch 1983.

ture of vague predicates – that feature of theirs marking the source of all their characteristic properties – simply is standard-sensitive context dependency.

The key ingredient of this argument is, of course, the theory of vague selections. After all, standard-sensitivity does not by itself imply the ability to have borderline cases in *all* contexts, nor the presence of higher-order vagueness, nor, as we shall see in the next chapter, the susceptibility to sorites paradoxes. Instead, standard-sensitive predicates take on all these properties precisely because vague selections are so integral to their interpretation, and vague selections happen to be saddled with a distinctive sort of "haziness". In this sense, linguistic vagueness is metaphysically reduced to the "vagueness" of these selections. By the end of the next chapter, we will have shown that the existence of borderline cases and the susceptibility to sorites paradoxes – perhaps the most recognized properties of vague predicates – are features of human language precisely because vague selections have all those characteristic uncertainty orders and are used in the interpretation of so many natural language expressions. In a real sense, then, vague selections are the source of all vagueness in language, and it is in this sense that our theory offers a metaphysical reduction of linguistic vagueness to the "vagueness" of vague selections.

Chapter 4

The Heap

Vague Selections and the Sorites

4.1 Chapter Overview

We at last come to that subject with which most studies of vague predicates begin: the sorites paradox. The reader is warned that this final chapter contains the most speculative and unpolished material of the thesis. One might say that the road is rough, but the weather is nice. Even the lightest stroll through this chapter will still jarringly lead the reader over a great many lacunae, yet he will be continually cooled by a soft, pleasant breeze emanating from some olympic style hand-waving.

The goal of this chapter is to demonstrate that our dynamic semantics — and in particular, the second Vagueness Principle — provides an original resolution to the sorites paradox. Furthermore, when this perspective on the sorites is combined with our theory of predication from Section 3.8.3, one may deduce that predicates are sorites-susceptible if and only if they are standard-sensitive.

The chapter begins in the next section with a defense of the fundamental assumptions regarding the sorites from which our analysis starts. Specifically, we defend the claims that the sorites argument does not reveal a genuine paradox, and that the proper response to the argument is to reject one of its premises, rather than to restrict one's reasoning from those premises.

In Section 4.3, we ready ourselves for a formal attack upon the sorites by our dynamic semantics. The notions of "sorites susceptibility" and "sorites series" are retooled to interconnect more amenably with the technology of our formal theory in Chapter 3. This train of thought then abruptly breaks off in the following sections, as we introduce two important components of our analysis. Section 4.4 addresses the peculiarities of asserting the existence of precise boundaries within well orderings, while Section 4.5 discusses and then precisely formulates the notion that vague selections have a "minimal size." This lower bound on the girth of vague selections is then argued to imply a law further constraining the structure of contexts, one that will have a great importance in our theory of the sorites.

Section 4.6 unveils our analysis of the sorites. In a soundbite: the size restriction on vague selections implies a minimum degree of difference between two objects which can be identified in a context as lying on opposite sides of the extension of a vague predicate. This is argued to provide a resolution to the paradox somewhat akin to that found in Graff 2000. Section 4.7 generalizes this analysis to the "higher order" versions of the sorites. Finally, in Section 4.8, we argue that a slightly revised version of the analysis may be combined with our theory of predication from Section 3.8 to predict that predicates are sorites susceptible if and only if they are standard-sensitive.

Then, the thesis is over.

4.2 Paradox, Fallacy, or Misunderstanding?

In this section, we defend our position upon the most fundamental questions concerning the sorites. Although in recent years, theorists have tended to leave their position on these matters unargued, and skip directly to the peculiarities of their analyses, we include this material because (i) we feel that it offers an original and helpful articulation of these points, and (ii) the constraint FINALITY outranks BREVITY in our OT tableau.

4.2.1 What is at Issue

What should one say about the sorites argument? Does it reveal a genuine paradox? Or, is it some sly trick of language we have yet to sort out properly? If the latter, how is the argument to be answered? Do we disagree with one of its premises, or do we maintain that the conclusion does not follow?

In considering these questions, let us take the following as our concrete specimen of the argument: a seven foot tall man is tall; a five foot tall man is not tall; any man only one inch shorter than a tall man is himself tall; thus, a man who is six foot eleven is tall, and so a man who is six foot ten is tall, and so a man who is six foot nine is tall... and so a man who is five feet is tall. Contradiction.

When it isn't introduced as a mere parlour game, this argument is usually intended to demonstrate that the meaning of the word "tall" is incoherent, that the rules governing the application of the term require a speaker to believe two contradicting propositions [1]. When faced with this argument, the most basic question one must address, then, is "Does it succeed?" Does this argument reveal a genuine inconsistency in our use of vague predicates? If so, then the sorites paradox would be an argument akin to the liar paradox or Russell's antinomy. Both these arguments are unanimously recognized as demonstrations that our intuitive conceptions of "truth" and "set" are too weak to be of use in precise theoretical reasoning; one can, by clever manipulation of their meaning, create

[1] See, for example, Dummett 1975.

a situation in which a speaker feels committed to asserting two contradicting statements.

If the sorites argument were considered successful, we would be immediately beset by the task of explaining how vague language may be so useful in spite of its basic incoherency. Since an overwhelming percentage of natural language is vague, it would seem a flat-out miracle that we aren't conspicuously drawn into contradiction on a daily basis. To draw a snide parallel: the words "fascism" and "democracy" in modern political discourse have arguably been divested of all coherent meaning [2]. There are totalitarian dictatorships defended as "democracies" and free and happy republics criticized as "fascist police states." To a dispassionate observer or alien linguist, there would appear no coherent set of rules governing the applications of these terms. However, the sort of blatant contradiction, hypocrisy and incomprehensibility we find accompanying these words does not seem to happen when people speak about colors, tallness or heaps of things. Why is vague language so robustly well behaved? Furthermore, as Veltman 2001 points out, we greatly prefer the use of vague expressions to that of precise ones; we find them more informative. If I were to say to you "My australian shepherd can run an Agility 3 course in 3 minutes, 32 seconds," you would most likely respond "So... is that fast?" I have already named a precise time span within which my dog can run the course, but you wish to know whether such an act may be described by the vague predicate "fast." If "fast" were an incoherent idea, then, strictly speaking, it would carry no information value. So, why does it feel vastly more informative than other, more precise descriptions?

Accepting the sorites argument as successful appears to bring some difficult challenges. Sometimes, of course, the truth hurts. But, is there much reason for accepting the sorites? An unschooled speaker is unable to refute the argument, but that needn't mean that the argument is valid or all its premises true. The argument may be incorrect, and our inability to say *why* it's incorrect the result of our not having a sophisticated enough reflective understanding of how our own vague concepts operate. An example of just such a "false paradox" would be Zeno's paradoxes concerning motion. These arguments attempt to demonstrate that our most natural conceptions of space and motion are incoherent. Although the average person cannot answer these arguments, it is nevertheless widely held that the field of differential calculus developed in the Eighteenth Century vindicates our common conceptions of space and motion, explains in greater detail how those concepts work, and thereby answers Zeno's arguments.

Now, let us look closely at the sorites argument. Which is the sorites most like: Russell's antinomy or Zeno's paradox? One would do better to lump it with Zeno's paradox.

[2]Indeed, if Orwell is to be believed, they were sapped of their sense well before the end of the second world war.

4.2.2 The Sorites is Not Successful

There are three facts supporting the falsity of the sorites. The first is that the argument sounds "funny" or "suspicious" to the untrained ear, and there is a fairly good track-record of errors being uncovered in "funny"-sounding would-be paradoxes. The second is that one feels no compulsion to believe the contradictory conclusion of the argument, which suggests that the argument does not, after all, conform to the rules governing vague language. The third is that there is simply no independent evidence that vague language is incoherent. Let us consider each in detail.

Circumstantial evidence in favor of our lumping the sorites with Zeno's paradox lies in how similar are our common-sense reactions to the two arguments. It takes a sophisticated intellect to believe the sorites, and this distinguishes it from the liar paradox while drawing it closer to Zeno's arguments. Just as no normal person has ever believed the sorites argument to show that our everyday vague concepts are incoherent [3], no one has earnestly believed Zeno's paradoxes to show that our concepts of space and time are incoherent [4]. To the average person, living in Alexandria as in New York, Zeno's paradox and the sorites sound "wrong"; they appear "funny." As with simpler false paradoxes, such as the paradox of the preface or the paradox of the surprise quiz, the average person detects that there is something "queer" about Zeno's reasoning, though he cannot quite put it into words or articulate a refutation. Indeed, when told the accepted refutation, his head may very well spin, but nevertheless, his implicit knowledge of the operation of his concepts is sufficient for him to recognize that such reasonings as appear in these arguments are illicit.

Now, compare these false paradoxes to the liar paradox or to Russell's antinomy. No one finds anything wrong with the argument that the liar's sentence must be both true and false if it is either, or with the argument that the Russell set must be both a member of itself and not a member of itself. A general distaste for admitting the existence of paradox, then, cannot be what is responsible for the widespread skepticism concerning the sorites. One suspects that, like Zeno's paradox, the paradox of the preface, and the paradox of the surprise quiz, our reluctance to accept the sorites — our feeling that there is something "wrong" with it — is due to our detecting something genuinely hazy within the reasoning, something which our semantic theory of vague predicates simply has yet to draw out into the light of day.

Of course, it could also be that the sorites argument sounds "funny" merely because its way of reasoning with vague predicates, though legitimate, is unfamiliar. We don't naturally employ such arguments in our daily life, and so the sorites strikes us as a novelty. There is, however, a more direct line of evidence that the sorites fails. Recall that the sorites beginning this section intends to show that the rules governing the predicate "tall" commit a speaker to two

[3] "Normal" here means, approximately, not a linguist or philosopher, and certainly not Michael Dummett or Peter Unger; see, for example, Unger 1979.

[4] Except for maybe Zeno, of course. But, Diogenes himself replied that "it is solved by walking" (solvitur ambulando).

contradictory statements. For the argument to succeed, it must construct from these rules a deduction that a five foot man is both tall and not tall. But, no speaker accepts the constructed deduction that a person five feet in height is tall. No one feels in the least that they must apply the adjective "tall" to Michael Jackson simply because there is a sorites series linking him to Michael Jordan. Indeed, despite the existence of that series, speakers feel *compelled* to describe Michael Jackson as "not tall" and to disagree with anyone who believes otherwise. However, if the meaning of "tall" sanctioned the deduction used in the sorites argument, people should recognize it as providing them reason for asserting that Michael Jackson is tall. That they don't demonstrates it doesn't.

Thus, the dirty fact that average speakers are never led by the sorites into the assertion that Michael Jackson is tall reveals that the argument does not truly conform to the principles governing vague predicates, contrary to its aims. Again, the sorites can be contrasted with the liar paradox. Everyone agrees with the reasoning behind the liar paradox; everyone agrees that it shows there to be equal reason for labeling the liar's sentence 'true' and for labeling it 'false'. Thus, the liar paradox, by creating a situation in which people honestly perceive themselves as committed to two contradicting propositions, succeeds in its aims, while the sorites, by not creating such a situation, fails.

Both of these points are compactly and acutely summarized in Wright 1987.

> "This brings us to the third, and, I think, a decisive objection to (the view that the sorites genuinely creates a paradox – SC). I do not see how we can rest content with the idea that certain implicitly known semantic rules are incoherent when *nobody's* reaction, on being presented with the purported demonstration of the inconsistency, i.e., the paradox – even if they can find no fault with it – is to lose confidence in the unique propriety of the response – e.g. "That's orange" – which the demonstration seems to confound. Think of your reaction when, having received as explanation of the notion of class only the usual informal patter plus the axioms of naive set theory, you first confronted Russell's paradox. If, which is unlikely, you held any intuitive conviction about whether Russell's class was a member of itself, you will have been forced to recognize an exact parity in the opposing case; and the effect should have been to cause you to realize that there is just no view to take on the matter before some refinement of the notion of class has been made. But that is exactly *not* our response to the sorties paradox for "red". Our conviction of the correctness of the non-inferential ingredient in the contradiction is left *totally undisturbed.* So far from bringing us to recognize that, pending some refinement in the meaning of "red", there is just no such thing as justifiably describing something as "red" or not, our conviction is that no one *ought* to be disturbed by the paradox – and this conviction is not based on certainty that we shall be able to disclose some simple fallacy. If the rules for the use of "red" really do sanction the paradox, why do we have absolutely no sense of dis-

turbance, no sense that a *real case* has been made for the inferential ingredient at all? ... A different account is called for." (Wright 1987)

Finally, as was already mentioned, the success of the sorites argument would imply that an activity which can largely be described via consistent sets of rules is nevertheless governed by rules which are inconsistent. It would be an incredible feat for our use of vague language to have the coherency it is observed to have, were the principles underlying it to have none. After all, the *only* situations in which this supposed incoherency is claimed to reveal itself are those in which sorites series are exhibited. Why doesn't the inconsistency become apparent *anywhere else*? Recall, too, that not even these situations reveal a genuine inconsistency of use, since speakers never accept the contradictory conclusion of the sorites argument. At best, the sorites reveals a situation in which it is impossible to justify one's consistent behavior, not in which one behaves inconsistently. So, aside from the fact that no one has yet convincingly refuted the sorites argument, there is no evidence to support the claim that vague language is governed by an inconsistent logic, and so the task of explaining how, despite its overwhelming appearance of consistency, the logic of vague predicates is inconsistent, looks uniquely unpromising.

4.2.3 The Sorites Rests on a False Premise

We conclude that it puts one in a far better position to maintain that the sorites is a false paradox than to accept it as genuine. Putting oneself in this position, however, carries its own hassles. If the sorites paradox is false, what, then, is its refutation? Is the reasoning unsound, or does it contain a false premise? There is a rough consensus in the literature that the best way to resolve the sorites paradox is by rejecting one of its premises, and nearly always the "inductive" premise of the argument is targeted for elimination.

Our choice of swimming with the intellectual currents on this point is both principled and opportunistic. It is opportunistic in that we have already demonstrated, in Chapter 2, that our semantic analysis of vague predicates renders sound all the classical principles of logic required to deduce from the sorites premises the contradictory conclusion. Therefore, unless we want to wreck the home we've made, we will have to accept that the four sorites premises are jointly inconsistent. Our choice is also principled in that the tradition of attacking the premises of the sorites is no mere fad or rut, but is based on convincing arguments made by a number of philosophers that any attempt to render all four premises of the sorites consistent will just not prove effective. We will not here rehearse those arguments in detail; instead, we refer the reader to the works of Cargile 1969, Dummett 1975, and Kamp 1981. In brief, the problem is that the contradictory conclusion may be derived from the sorites premises by means of very minimal logical assumptions, ones which are accepted even by constructivists and strict finitists. Individually, each of these assumptions seems quite obviously true, and their rejection to fly in the face of clear linguistic data. Therefore, frightened off from this path, we will consider the possibility of rejecting one of the argument's premises.

When we look to the premises, the inductive premise jumps out immediately. What, for example, could be more obvious than the fact that Michael Jordan is tall and Michael Jackson isn't. Moreover, is anyone prepared to deny that a series of men, each only one inch taller than the other, could be created, one which begins and ends with our two Michaels? The inductive premise, on the other hand, sounds a bit "funny." "If one man is tall, then a man only one inch shorter than him is also tall." What an odd thing to say. I am not so confident. When I picture the meaning of this sentence in my head, I picture the ridiculousness of there being before me two men, only one inch apart in their height, yet I say that one man is tall and the other is not tall. But, simply because that scenario is weird, just because it could never happen, does that entail that the inductive premise *as it's stated above* is true? Common sense screams at me "Of course, you overeducated egg-head!", but there seems to be some room to wriggle here, and in the following sections, we will do a good bit of wriggling.

4.2.4 What Now Must be Done

Although we may have settled into what appears the most comfortable position, we are never fully relieved of discomfiting tensions. In giving up the truth of the inductive premise, one must still answer a number of unpleasant questions. If we deny that "If one man is tall, then a man only one inch shorter than him is also tall", and we have retained all of classical logic, then the statement "There exists a man who is tall and a man only one inch shorter than him who isn't" must be true. Why, then, does this seem so ridiculous? Why does this sentence appear straightforwardly false?

In Graff 2000, the task now before us is given a clear and helpful formulation. Graff 2000 identifies two questions which any analysis of the sorites that rejects its inductive premise must answer: "The Epistemological Question" and "The Psychological Question." The Epistemological Question asks, if the inductive premise is not true, why can one not find any falsifying instances? Informally, why does it strike me as impossible for there to be two men, only an inch difference in their height, one of whom is tall and the other not tall? The Psychological Question asks, if the inductive premise isn't true, why does common sense compel us to accept it? Notice that these questions are distinct; one may well imagine an analysis of the sorites that explains how we could believe the inductive premise without explaining why falsifying instantiations for the premise cannot be found.

In the remainder of this chapter, we will show that our theory of vague selections helps us towards answers to both the epistemological and psychological questions.

4.3 Approaching the Sorites

One of our goals in this chapter is to deduce, by means of our semantic analysis of standard-sensitive predicates, that all standard-sensitive context dependent predicates are susceptible to sorites paradoxes. Therefore, we will need to reformulate the notion of "sorites susceptibility," using the concepts made available within our formal semantics, so that such a deduction within our formal theory might be possible. For this task, we will first propose a definition which comes rather close to our more informal and concrete notion of "sorites susceptibility" presented in Section 1.1. This conservative reformulation, we will then extend into greater abstraction.

Concept 3.1 (Sorites Series I) Let P be a standard-sensitive predicate, O^P be its densely ordered set of degrees, and $\langle K, \langle w, \delta \rangle\rangle$, $K = (W', R', V'), W' \subseteq W \times O^P$ be a context. A **sorites series** for P at $\langle K, \langle w, \delta \rangle\rangle$ is a subset S of the domain of objects at $\langle K, \langle w, \delta \rangle\rangle$ [5], well-ordered under a relation Q such that Q has a first element α, a last element β, $\langle K, \langle w, \delta \rangle\rangle \Vdash DP\alpha$, $\langle K, \langle w, \delta \rangle\rangle \Vdash D\neg P\beta$ and for all $\gamma, \eta \in S$, if $Q\gamma\eta$, then it *appears* [6] to the speakers in the context that $\langle K, \langle w, \delta \rangle\rangle \Vdash D(P\gamma \rightarrow P\eta)$.

Definition 3.2 (Sorites Paradox) Let P be a standard-sensitive predicate, O^P be its densely ordered set of degrees, and $\langle K, \langle w, \delta \rangle\rangle$, $K = (W', R', V'), W' \subseteq W \times O^P$ be a context. There exists a **sorites paradox** for P at $\langle K, \langle w, \delta \rangle\rangle$ if there exists a sorites series for P at $\langle K, \langle w, \delta \rangle\rangle$.

Definition 3.3 (Sorites Susceptible) Let P be a standard-sensitive predicate, and O^P be its densely ordered set of degrees. P is **sorites susceptible** if there is some context $\langle K, \langle w, \delta \rangle\rangle$, $K = (W', R', V'), W' \subseteq W \times O^P$ at which there is a sorites paradox for P.

We hope that the reader finds these definitions of "sorites series," "sorites paradox" and "sorites susceptible" intuitively acceptable, and a natural formalization of the notions appearing in Section 1.1. Definition 3.1 states, for

[5] By "the domain of objects at $\langle K, \langle w, \delta \rangle\rangle$", we mean the domain of objects at w, which may include more than simply those entities which exist in w.

[6] The notion of a context's "appearance" to its speakers will be left unformulated here, this constituting the grossest of the lacunae mentioned in the Chapter Overview and the sole reason why Concept 3.1 isn't named "Definition 3.1" (since it isn't, strictly speaking, a correct definition). Ideally, I would like it that it *appears* to the speakers in context $\langle K, \langle w, \delta \rangle\rangle$ that $\langle K, \langle w, \delta \rangle\rangle \Vdash \Phi$ iff $\langle K, \langle w, \delta \rangle\rangle \Vdash D\Phi$. As will become clear later, I would therefore like my logic for D to be such that $\langle K, \langle w, \delta \rangle\rangle \Vdash D\forall x, y((DPx \wedge Qxy) \rightarrow \neg D\neg Py)$ implies $\langle K, \langle w, \delta \rangle\rangle \Vdash DD\forall x, y((Px \wedge Qxy) \rightarrow Py)$. The logic I have now obviously doesn't permit me this, and I can think of no easy augmentations of it which do. On the other hand, this concept of "appearance" would also seem to require that for any such context $\langle K, \langle w, \delta \rangle\rangle$, if $\langle K, \langle w, \delta \rangle\rangle \Vdash D\phi$, then it appears to the speakers in $\langle K, \langle w, \delta \rangle\rangle$ that $\langle K, \langle w, \delta \rangle\rangle \Vdash D\phi$.

Thus, this concept of appearance is not presently in a very happy state.

example, that a sorites series for a predicate P at a context is simply a well-ordered series of objects taken from the domain available at the context, having a first and last member such that the speakers in the context may assert that the first member is P and that the last member is $\neg P$, and it appears to the speakers that, for any two objects directly adjacent in the series, they may assert that if the one is P then the other is P as well. This certainly feels to capture the essential features of a sorites series.

The reader, however, should not become attached to these definitions; they are merely a stepping stool. For the purposes of achieving a greater degree of abstraction, of simplifying our discussion, and of easing the length of our proofs, we will discard Concept 3.1 and use instead a more abstract formal representation of sorites series. The intuition behind this more abstract definition is that, in all actual sorites series for a vague predicate P, the ordering of the objects making up the series is isomorphic to the ordering of their degrees of P. In our newer definition, then, we "strip away" the concrete objects making up the sorites series, and leave only the ordering of their degrees in place.

First, however, we must introduce some new notation. If $S \subseteq O^P$, then let $\prec_p^s \subset S \times S$ be the relation $\{\langle \alpha, \beta \rangle : \alpha <_P \beta$ and there is no $\gamma \in S$ such that $\alpha <_P \gamma <_P \beta\}$. Furthermore, we add "$\prec_p^s$" as a symbol within our object language, giving it the obvious semantic definition that $\langle K, \langle w, \delta \rangle \rangle \Vdash \alpha \prec_p^s \beta$ iff $\alpha \prec_p^s \beta$.

Concept 3.4 (Sorites Series II) Let P be a standard-sensitive predicate, O^P be its densely ordered set of degrees, and $\langle K, \langle w, \delta \rangle \rangle, K = (W', R', V'), W' \subseteq W \times O^P$ be a context. A **sorites series** for P at $\langle K, \langle w, \delta \rangle \rangle$ is a subset S of O^P, such that $\prec_p^s$ is a well-ordering with a first element α and a last element β[7], $\langle K, \langle w, \delta \rangle \rangle \Vdash DP\alpha$, $\langle K, \langle w, \delta \rangle \rangle \Vdash D\neg P\beta$ and it *appears* to the speakers in $\langle K, \langle w, \delta \rangle \rangle$ that $\langle K, \langle w, \delta \rangle \rangle \Vdash D\forall x, y((Px \wedge y \prec_p^s x) \rightarrow Py)$.

Throughout the rest of this chapter, we will use the term "sorites series" with the meaning supplied by Concept 3.4. Moreover, Definitions 3.2 and 3.3 are reinterpreted in light of this new definition.

It should be clear to the reader that at any context at which there is for a predicate P a sorites series by Concept 3.4, and at which there are objects in the domain witnessing each of the degrees in that sorites series, there will be a sorites series for P by Concept 3.1. Therefore, we maintain that Concept 3.4 handily formalizes the notion of there being at a context an *imaginable* sorites for P. The reader may observe the resemblance here to our formulation in Section 3.7.3 of the notion of there being an imaginable borderline case for P.

Finally, to further reduce the size of our arguments, we will introduce here the concept of a "constraining selection."

[7]We violate the standard conventions here, and use the term "first element" to mean an $\alpha \in S$ such that for all $\gamma \in S$, where $\gamma \neq \alpha$ and $\prec_p^{sT}$ is the transitive closure of $\prec_p^s$, $\gamma \prec_p^{sT} \alpha$. A similarly heretical definition is assumed for "last element."

Definition 3.5 Let $\langle K, \langle w, \delta \rangle \rangle$, $K = (W', R', V')$, $W' \subseteq W \times O^P$ be a context, and $\langle K', w \rangle \in$ **vagsel**, $K' = (W, R, V)$. $\langle K', w \rangle$ is a **constraining selection for** $\langle K, \langle w, \delta \rangle \rangle$ if for all $\langle w, \delta \rangle \in W'$, $\{w' : \exists \delta' \in O^P\ R' \langle w, \delta \rangle \langle w', \delta' \rangle\} = \{w' : Rww'\}$ and $\{\delta' : \exists w' \in W\ R' \langle w, \delta \rangle \langle w', \delta' \rangle\} = \{\delta' : \langle K', w \rangle \Vdash sel(\delta')\}$.

By the second Vagueness Principle (Principle 3.7.12), there exists a constraining selection for every context.

Now, it is easily shown that, by our dynamic semantics, the four premises of the sorites argument are jointly inconsistent.

Proposition 3.6 There is no context $\langle K, \langle w, \delta \rangle \rangle$ such that there exists a well-ordering Q with a first element α and a last element β, $\langle K, \langle w, \delta \rangle \rangle \Vdash DP\alpha$, $\langle K, \langle w, \delta \rangle \rangle \Vdash D\neg P\beta$ and $\langle K, \langle w, \delta \rangle \rangle \Vdash D\forall x, y((Px \wedge Qxy) \rightarrow Py)$.

PROOF: Suppose there were such a context $\langle K, \langle w, \delta \rangle \rangle$. The second Vagueness Principle implies the existence of a constraining selection for $\langle K, \langle w, \delta \rangle \rangle$. Now, by the semantics of D and the existence of a constraining selection, there is some $\langle w', \delta' \rangle$ such that $R' \langle w, \delta \rangle \langle w', \delta' \rangle$ and $\langle K, \langle w', \delta' \rangle \rangle \Vdash P\alpha$, $\langle K, \langle w', \delta' \rangle \rangle \Vdash \neg P\beta$ and $\langle K, \langle w', \delta' \rangle \rangle \Vdash \forall x, y((Px \wedge Qxy) \rightarrow Py)$. Induction on the elements of the ordering Q is now sufficient to prove that $\langle K, \langle w', \delta' \rangle \rangle \Vdash P\beta$, contrary to hypothesis. Note that this argument is indifferent as to whether the elements of Q are degrees in O^P or objects in the domain of $\langle K, \langle w, \delta \rangle \rangle$.

Just as we had adumbrated, then, our dynamic semantics implies that one of the premises of the sorites argument must not be true. Recalling Graff 2000's psychological question, the task now before our semantic theory is to explain why all four premises of the argument should appear to be simultaneously true. That is, we must explain why it seems to us as speakers that there are contexts which Proposition 3.6 demonstrates not to exist.

Now, by either Concept 3.1 or 3.4, a context in which there is a sorites series for a standard-sensitive predicate P will be one that appears to its speakers to support all the premises of the sorites argument. Thus, if we could somehow prove that for any standard-sensitive predicate P and context $\langle K, \langle w, \delta \rangle \rangle$, there exists a sorites series for P at $\langle K, \langle w, \delta \rangle \rangle$, we could accomplish two goals of this chapter at one stroke. First, by Definition 3.3, we will have shown that all standard-sensitive predicates are sorites susceptible [8]. Secondly, our hypothetical proof would need to show that, for any context $\langle K, \langle w, \delta \rangle \rangle$, it appears to speakers at $\langle K, \langle w, \delta \rangle \rangle$ that $\langle K, \langle w, \delta \rangle \rangle \Vdash D\forall x, y((Px \wedge y \prec_p^s x) \rightarrow Py)$. Such a proof would, in effect, have to explain why speakers find the inductive premise of the sorites argument so darn plausible, and so would answer Graff 2000's psychological question.

Our goal, then, is some argument to the effect that for any standard-sensitive

[8] Although this point will not be pursued here, we will have also shown, by our comments above regarding Concept 3.4, that for any standard-sensitive predicate P, there exists in any context an *imaginable* sorites series for P.

predicate P and any context $\langle K, \langle w, \delta \rangle \rangle$, a sorites series $\langle S, \prec_p^s \rangle$ for P at $\langle K, \langle w, \delta \rangle \rangle$ may be constructed. The crafting of this argument, however, will require the composition of some rather disparate facts; Sections 3.4 and 3.5 bring those facts to light.

4.4 Precise Boundaries Within Well Orderings

I am the captain of a basketball team, and I must choose which players will be the first ones in the game and which will sit out at the beginning. Let us imagine that the players are arranged in a straight line, and I am to divide that line cleanly between the starting players and the reserves [9]. What I must do in such a task is *assert* of one of the players in the ordering that he will sit out and *assert* of his neighbor in the ordering that he will be in the starting lineup. Drawing this boundary between the players makes it common, definite knowledge which are the starting players and which are not. Moreover, for any two players adjacent in the ordering, this act of boundary drawing makes it common, definite knowledge whether either one is a starting player.

This relationship between boundary drawing within well-orderings, assertion and definite knowledge appears rather robust. Imagine that I am a security guard backstage of a rock concert, and I am presented with a list of names of persons who want to come to meet the band. The list is obviously too long; the dressing rooms can't hold 500 people. Therefore, what I must do is assert of one of the names in the list that it is the "last" to be allowed backstage. Doing so indirectly asserts of an adjacent name in the list that it is the "first" not to be allowed backstage. My act of drawing this line also makes it common, definite knowledge of every two names adjacent in the list, whether the persons bearing those names are to be permitted backstage.

Let us, then, jump to the conclusion that any act of drawing a precise boundary for some property within a well-ordering entails that there are at least two entities directly adjacent in the ordering such that it is definite knowledge whether the property holds of them. This is stated more precisely in the following law.

Principle 4.1 (Law of Precise Boundaries) Let $\langle S, \leq \rangle$ be a well ordering, P be a predicate, and $\langle K, \langle w, \delta \rangle \rangle$ be a context. Drawing a precise boundary within $\langle S, \leq \rangle$ at $\langle K, \langle w, \delta \rangle \rangle$ for P entails the existence of two objects $\alpha, \beta \in S$ such that $\langle K, \langle w, \delta \rangle \rangle \Vdash DP\alpha$, $\langle K, \langle w, \delta \rangle \rangle \Vdash D\neg P\beta$, $\alpha \leq \beta$ and for no $\gamma \in S$ such that $\gamma \neq \alpha \neq \beta$ does $\alpha \leq \gamma \leq \beta$.

[9] This is, of course, not the best way to manage one's team, but it does develop my point.

4.5 The Sizes of Vague Selections

4.5.1 The Minimality of Vague Selections

At one point in Section 3.4, vague selections were described as limited in their "size," in that one could not with a vague selection attain an arbitrary degree of precision. This observation eventually lead to our introduction of the Selection Principle. In this section, we will dwell on the observation a bit longer, and propose a more precise, formal statement of it, one which will have a great importance later on.

Often when I make a vague selection, it seems that I could have made one which was, in some sense, "choosier" or "more precise." Suppose, for example, that you and I are designing a set for a play. We are standing in front of a raised piece of plywood. I ask you to cut out a square from the plywood "there," and wave my hands towards some section of the board. Nervous about not meeting my desires, you respond "Where?", thereby eliciting from me a more precise indication of area. "Here," I say, and pat my hand on a particular portion of the board. Still somewhat apprehensive, you ask, "Wait, show me where you want it." Somewhat frustrated, I snap back "Ach! Right here!" and trace with my finger a square on the plywood. By the criteria laid out in Section 3.4, all of my imagined actions were vague selections, but in some sense each was more "precise," or "smaller," than that before it.

There seems to be a limit to this process, however. You may continue to question me, but there comes a point at which, without the use of technical instruments or other artificial means, I simply cannot continue to sharpen my intentions. I may squint harder, I may move closer to the board and scour its surface more slowly, but there is an intuition that at some point I simply do not achieve any greater precision in either what I communicate or what I myself even intend. This point was also made in Section 3.4 regarding one's ability to select, or attend to, a point in space. I can, of course, make vague selections of spatial points to varying degrees of precision. There is nevertheless a limit to the precision my thoughts can take; they plateau at some point, and I can no longer imagine points any "smaller."

Let us accept, then, that vague selections have a minimum "size." We might state this more precisely through the use of a measure function μ. In order to do this, let us for the moment drop our earlier formalization of vague selections, and instead think of them momentarily as collections of selected points. Where s is a vague selection of points drawn from a dense linear ordering $\langle O^P, <_P \rangle$, sel is a predicate ranging over the elements of O^P and denoting those elements selected by s, and I is our indeterminacy operator, we may represent s as simply the set $\{\alpha \in O^P : \exists n \in \mathbb{N}\ I^n sel(\alpha)\}$. The vague selection s, then, will be for the moment simply all those points for which it is to some degree indeterminate whether they have been selected. Let us also fix a measure function μ over the subsets of O^P. This function allows us a means of speaking of the "size" of various subintervals of O^P and the "distance" between points in the ordering $\langle O^P, <_P \rangle$. We may now state the existence of a minimum "size" for the vague selections over O^P in the following way.

Principle 5.1 (Minimality of Vague Selections) Let s be a vague selection of points drawn from a dense linear ordering $\langle O^P, <_P\rangle$, sel be a predicate ranging over the elements of O^P and denoting those elements selected by s, and μ be the fixed measure function over the subsets of O^P. There is a real number $\epsilon \in \mathbb{R}$ such that $\mu(\{\alpha \in O^P : \exists n \in \mathbb{N}\ I^n sel(\alpha)\}) \geq \epsilon$.

4.5.2 The Minimality of the Uncertainty Orders

In our construing the vague selection s as simply the set $\{\alpha \in O^P : \exists n \in \mathbb{N}\ I^n sel(\alpha)\}$, we thereby think of the vague selection as being composed of its various "uncertainty orders." Moreover, its not unreasonable to suppose that as the "size" of the vague selection changes, the various uncertainty orders constituting it stay "in proportion" to one another, like the pieces of a JPEG manipulated in Photoshop. It does seem upon reflection that as the vague selection becomes "larger," as the intentions behind the selection become looser, the span both of selected and indeterminately selected points increases, and that there is a fixed relationship between the sizes of these spans.

If we accept the ideas that vague selections are composed of their uncertainty orders and that these orders have a fixed, proportional relationship, we may derive from Principle 5.1 that there is also a minimal "size" to each of the uncertainty orders within a vague selection.

Principle 5.2 (Minimality of Uncertainty Orders) Let s be a vague selection of points drawn from a dense linear ordering $\langle O^P, <_P\rangle$, sel be a predicate ranging over the elements of O^P and denoting those elements selected by s, and μ be the fixed measure function over the subsets of O^P. There exists a function $f : \mathbb{N} \to \mathbb{R}$ such that for all n, $\mu(\{\alpha \in O^P : I^n sel(\alpha)\}) \geq f(n)$.

ARGUMENT: The following is not intended as an exact proof of Principle 5.2 [10], but more as an informal argument intended to spell out the intuition behind it. It will introduce some facts and machinery that will be of help along the way.

We may state the fact that vague selections are "composed" of their various uncertainty orders as $\{\alpha \in O^P : \exists n \in \mathbb{N}\ I^n sel(\alpha)\} = \bigcup_{n=0}^{\omega}\{\alpha \in O^P : I^n sel(\alpha)\}$. Thus, $\mu(\{\alpha \in O^P : \exists n \in \mathbb{N}\ I^n sel(\alpha)\}) = \mu(\bigcup_{n=0}^{\omega}\{\alpha \in O^P : I^n sel(\alpha)\}) = \Sigma_{n=0}^{\omega}\mu(\{\alpha \in O^P : I^n sel(\alpha)\})$. By the Minimality of Vague Selections, then, there is an $\epsilon \in \mathbb{R}$ such that $\Sigma_{n=0}^{\omega}\mu(\{\alpha \in O^P : I^n sel(\alpha)\}) \geq \epsilon$.

Now, since the various uncertainty orders have a fixed proportion, there is for each $n \in \mathbb{N}$ some $\delta \in \mathbb{R}$ such that $\mu(\{\alpha \in O^P : I^n sel(\alpha)\}) = \delta$ whenever $\Sigma_{n=0}^{\omega}\mu(\{\alpha \in O^P : I^n sel(\alpha)\}) = \epsilon$ [11]. Let, then, $f : \mathbb{N} \to \mathbb{R}$ be a function which takes n to this real number δ. That is, $f(n) = \delta$ if $\mu(\{\alpha \in O^P : I^n sel(\alpha)\}) = \delta$ whenever $\Sigma_{n=0}^{\omega}\mu(\{\alpha \in O^P : I^n sel(\alpha)\}) = \epsilon$. Clearly, then, $\epsilon = \Sigma_{n=0}^{\omega} f(n)$.

[10] Otherwise, it would named "Theorem 5.2".

[11] This statement seems obviously true, but I know of no formal way of proving it.

Now, suppose that there is a vague selection s over the ordering $\langle O^P, <_P \rangle$, and there is some $n \in \mathbb{N}$ such that $\mu(\{\alpha \in O^P : I^n sel(\alpha)\}) < f(n)$. Since all uncertainty orders have a fixed proportion, we derive that for all $m \in \mathbb{N}$, $\mu(\{\alpha \in O^P\colon I^m sel(\alpha)\}) < f(m)$ [12]. Thus, $\Sigma_{n=0}^{\omega}\mu(\{\alpha \in O^P : I^n sel(\alpha)\}) < \Sigma_{n=0}^{n} f(n)$, and so $\Sigma_{n=0}^{\omega}\mu(\{\alpha \in O^P : I^n sel(\alpha)\}) < \epsilon$, contrary to the fact that ϵ is our lower bound. Therefore, for all n, $\mu(\{\alpha \in O^P : I^n sel(\alpha)\}) \geq f(n)$.

The Minimality of Uncertainty Orders entails that for any vague selection, there is a minimal size which the set $\{\alpha \in O^P : sel(\alpha)\} = \{\alpha \in O^P : I^0 sel(\alpha)\}$ can have. That is, in any vague selection, we know that $\mu(\{\alpha \in O^P : sel(\alpha)\}) \geq f(0)$. This fact will have an immense importance to our analysis of the sorites paradox.

To get a rough idea of where we're heading with all of this, observe that the second Vagueness Principle implies that for any context $\langle K, \langle w, \delta \rangle \rangle$ and vague predicate P, each n-th level vagueness order for P at $\langle K, \langle w, \delta \rangle \rangle$ corresponds to an $(n-1)$-th level uncertainty order for some vague selection over the degrees of P.

Proposition 5.3 Let P be a standard-sensitive predicate, O^P be its densely ordered set of degrees, $\langle K, \langle w, \delta \rangle \rangle$, $K = (W', R', V')$, $W' \subseteq W \times O^P$ be a context, and $\langle K', w \rangle$ be a constraining selection for $\langle K, \langle w, \delta \rangle \rangle$. For all $n \in \mathbb{N}$, $\langle K, \langle w, \delta \rangle \rangle \Vdash I^{n+1}P(\alpha)$ iff $\langle K', w \rangle \Vdash I^n sel(\alpha)$.

PROOF: By Lemma 3.8.2, for all $n \geq 1$, $\langle K, \langle w, \delta \rangle \rangle \Vdash I^{n+1}P(\alpha)$ iff $\langle K', w \rangle \in \alpha$-n-unc iff, by Lemma 3.7.5, $\langle K', w \rangle \Vdash I^n sel(\alpha)$. The case for when $n = 0$ follows from the admissibility of $\langle K', w \rangle$, the fact that it's a constraining selection for $\langle K, \langle w, \delta \rangle \rangle$, and the semantics for I.

In a picture, we might draw Proposition 5.3 as the following:

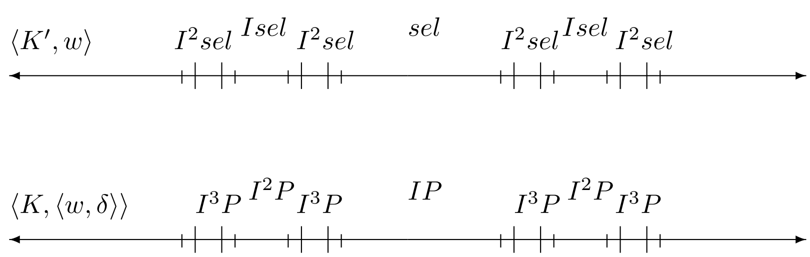

The idea which will be pursued in the following sections springs from the observation that the minimum size placed on the uncertainty orders of $\langle K', w \rangle$ by Principle 5.2 thereby places a minimum size on the vagueness orders for P

[12] Again, this statement seems obvious, but I know no way to precisely demonstrate it.

at $\langle K, \langle w, \delta\rangle\rangle$. This minimum size will then be used to construct a minimum "threshold" of difference between two degrees α and β such that it may be asserted in a context that $P\alpha$ and $\neg P\beta$. An answer to the epistemological and psychological questions may then be framed.

4.5.3 The Law of Quasi-Tolerance

Our discussion of vague selections and their "sizes" has thus far been relatively informal, and has not explicitly made use of the elements of the class **vagsel**. Let us now correct that. Moreover, let us import into our definition of **vagsel** the fact for any vague selection, there exists a lower bound on the size of the set of selected points. We shall do this by augmenting the original definition of an "admissible" pointed Kripke model.

Definition 5.4 Let $\langle O^P, <_P\rangle$ be our fixed dense ordering of degrees, μ be a fixed measure function over the subsets of O^P, and f be a fixed function $f : \mathbb{N} \to \mathbb{R}$. An **admissible** pointed Kripke model $\langle K, w\rangle$ is one such that if $K = (W, R, V)$, then for all $u \in W$, there exists an open, convex non-trivial subset S of O^P, with proper upper and lower bounds in O^P, such that for all $\alpha \in O^P$, $\langle K, u\rangle \Vdash sel(\alpha)$ iff $\alpha \in S$, and $\mu(\{\alpha \in O^P : \langle K, u\rangle \Vdash sel(\alpha)\}) \geq f(0)$.

The reader is invited to check that this change in our definition does not affect the results obtained in Chapter 3. Moreover, Definition 5.4 now allows us to quite easily prove the most important result of this section: the Law of Quasi-Tolerance. Let us first, however, explicitly introduce some standard notation. In the following, if $\alpha, \beta \in O^P$, we let $[\alpha, \beta]$ be the standard shorthand for the set $\{\delta : \alpha \leq_P \delta \leq_P \beta\}$.

Theorem 5.5 (Law of Quasi-Tolerance) Let P be a standard-sensitive predicate, O^P be its densely ordered set of degrees, $\alpha, \beta \in O^P$, and $\langle K, \langle w, \delta\rangle\rangle$, $K = (W', R', V')$, $W' \subseteq W \times O^P$ be a context. If $\langle K, \langle w, \delta\rangle\rangle \Vdash DP(\alpha)$ and $\langle K, \langle w, \delta\rangle\rangle \Vdash D\neg P(\beta)$, then $\mu([\beta, \alpha]) \geq f(0)$.

PROOF: Suppose that $\langle K, \langle w, \delta\rangle\rangle \Vdash DP(\alpha)$ and $\langle K, \langle w, \delta\rangle\rangle \Vdash D\neg P(\beta)$. By Principle 3.7.12, there exists a constraining selection $\langle K', w\rangle \in$ **vagsel** for $\langle K, \langle w, \delta\rangle\rangle$. By definition of **vagsel**, $\langle K', w\rangle$ is, of course, admissible.

Now, by the semantics of D, for all $\langle w', \delta'\rangle \in W'$ such that $R'\langle w, \delta\rangle\langle w', \delta'\rangle$, we have $\langle K, \langle w', \delta'\rangle\rangle \Vdash P(\alpha)$. Thus, by Definition 3.7.13, for every such $\langle w', \delta'\rangle$, $\alpha >_P \delta'$. Thus, α is a proper upper bound for the set $\{\delta' : \exists w' \in W\ R'\langle w, \delta\rangle\langle w', \delta'\rangle\} = \{\delta' : \langle K', w\rangle \Vdash sel(\delta')\}$. Similarly, from the fact that $\langle K, \langle w, \delta\rangle\rangle \Vdash D\neg P(\alpha)$, we may prove that β is a lower bound for the set $\{\delta' : \langle K', w\rangle \Vdash sel(\delta')\}$. Therefore, $[\beta, \alpha] \supset \{\delta' : \langle K', w\rangle \Vdash sel(\delta')\}$.

Now, suppose that $\mu([\beta, \alpha]) < f(0)$. By the principles of measure functions, then, $\mu([\beta, \alpha]) > \mu(\{\delta' : \langle K', w\rangle \Vdash sel(\delta')\})$, and thus $\mu(\{\delta' : \langle K', w\rangle \Vdash sel(\delta')\}) < f(0)$, contradicting $\langle K', w\rangle$'s admissibility. Thus, $\mu([\beta, \alpha]) \geq f(0)$.

Readers who are familiar with the standard literature on the sorites paradox and vague predicates will recognize that Theorem 5.5 is where our present analysis connects with the idea, often discussed in that literature, that vague predicates are "tolerant" of slight differences in the appearances of objects, and so are governed by a principle of "the equality of indiscernibles". Such readers will also note that Theorem 5.5 is a significantly weaker statement than the formulation which that principle often receives. Since the usual formulation of this principle of "tolerance" entails that the inductive premise of the sorites argument is true, our present perspective requires us to reject the claim that vague predicates are governed by that principle. Our discussion in Section 4.6.1, however, will reveal how the Law of Quasi-Tolerance above may be a more accurate description of the intuition supporting the usual formulation of this suspect principle.

Theorem 5.5 completes our set-up. In the next section, we finally bring together these various Principles, Propositions and Theorems into a coherent story.

4.6 Quasi-Tolerance and the Sorites Paradox

Our goal is to understand why it appears to us as speakers that there are contexts of the kind demonstrated by Proposition 3.6 to be impossible. That is, we want to understand why the inductive premise of the sorites argument always appears to us true, even though it never is. It was pointed out that this goal could be accomplished by demonstrating that for any standard-sensitive predicate P and any context $\langle K, \langle w, \delta \rangle \rangle$, a sorites series $\langle S, \prec_p^s \rangle$ for P at $\langle K, \langle w, \delta \rangle \rangle$ may be constructed. In this section, we will attempt that demonstration, and we begin with the following definition.

Definition 6.1 (Canonical Sorites Series) Let P be a standard-sensitive predicate, O^P be its densely ordered set of degrees, $\langle K, \langle w, \delta \rangle \rangle$, $K = (W', R', V')$, $W' \subseteq W \times O^P$ be a context, $\langle K', w \rangle \in$ **vagsel** be a constraining selection for $\langle K, \langle w, \delta \rangle \rangle$, μ and f be the fixed functions from Definition 5.4, and $\langle S, \prec_p^s \rangle$ be a well-ordering of degrees from O^P with a first element α and a last element β. $\langle S, \prec_p^s \rangle$ is a **canonical sorites series for** P **at** $\langle K, \langle w, \delta \rangle \rangle$, if α is a proper upper bound and β is a lower bound for the set $\{\delta' : \langle K', w \rangle \Vdash sel(\delta')\}$, and for any two $\gamma, \eta \in S$, if $\gamma \prec_o^s \eta$, then $\mu([\gamma, \eta]) < f(0)$.

Now we want to prove some facts about canonical sorites series, in particular that they are sorites series by Concept 3.4. First of all, it should be clear that for any context $\langle K, \langle w, \delta \rangle \rangle$ and standard-sensitive predicate P there is at least one canonical sorites series for P at $\langle K, \langle w, \delta \rangle \rangle$.

Proposition 6.2 Let Let P be a standard-sensitive predicate, O^P be its densely ordered set of degrees, $\langle K, \langle w, \delta \rangle \rangle$, $K = (W', R', V')$, $W' \subseteq W \times O^P$ be a context, $\langle K', w \rangle \in$ **vagsel** be a constraining selection for $\langle K, \langle w, \delta \rangle \rangle$, and

$\langle S, \prec_p^s \rangle$ be a canonical sorites series for P at $\langle K, \langle w, \delta \rangle\rangle$ with a first element α and a last element β. $\langle K, \langle w, \delta \rangle\rangle \Vdash DP\alpha$.

PROOF: Let $\langle w', \delta' \rangle$ be such that $R'\langle w, \delta \rangle\langle w', \delta' \rangle$. Thus, $\delta' \in \{\delta' : \langle K', w \rangle \Vdash sel(\delta')\}$. Now, since α is a proper upper bound for $\{\delta' : \langle K', w \rangle \Vdash sel(\delta')\}$, $\alpha >_P \delta'$. Thus, by Definition 3.7.13, $\langle K, \langle w', \delta' \rangle\rangle \Vdash P\alpha$. Since $\langle w', \delta' \rangle$ was arbitrary, $\langle K, \langle w, \delta \rangle\rangle \Vdash DP\alpha$.

Proposition 6.3 Let Let P be a standard-sensitive predicate, O^P be its densely ordered set of degrees, $\langle K, \langle w, \delta \rangle\rangle$, $K = (W', R', V')$, $W' \subseteq W \times O^P$ be a context, $\langle K', w \rangle \in$ **vagsel** be a constraining selection for $\langle K, \langle w, \delta \rangle\rangle$, and $\langle S, \prec_p^s \rangle$ be a canonical sorites series for P at $\langle K, \langle w, \delta \rangle\rangle$ with a first element α and a last element β. $\langle K, \langle w, \delta \rangle\rangle \Vdash D\neg P\beta$.

PROOF: Proved analogously to Proposition 6.2, only using the fact that β is a lower bound for $\{\delta' : \langle K', w \rangle \Vdash sel(\delta')\}$.

Proposition 6.4 Let P be a standard-sensitive predicate, O^P be its densely ordered set of degrees, $\langle K, \langle w, \delta \rangle\rangle$, $K = (W', R', V')$, $W' \subseteq W \times O^P$ be a context, $\langle K', w \rangle \in$ **vagsel** be a constraining selection for $\langle K, \langle w, \delta \rangle\rangle$, and $\langle S, \prec_p^s \rangle$ be a canonical sorites series for P at $\langle K, \langle w, \delta \rangle\rangle$ with a first element α and a last element β. $\langle K, \langle w, \delta \rangle\rangle \not\Vdash D\forall x, y((Px \wedge y \prec_p^s x) \rightarrow Py)$.

PROOF: This follows immediately from Propositions 3.6, 6.3 and 6.4.

We are now very close to establishing that canonical sorites series are sorites series by our definition in 3.4. Propositions 6.2 and 6.3, when combined with the peculiarities of Definition 6.1, reveal that canonical sorites series have nearly all the essential features of a sorites series. However, the most important detail is still missing: how does one argue that, despite Proposition 6.4, it should appear to speakers in a context that the context supports the untrue inductive premise $\forall x, y((Px \wedge y \prec_p^s x) \rightarrow Py)$ for the canonical sorites series $\langle S, \prec_p^s \rangle$? To answer this will require an answer to Graff 2000's psychological question, which we will have on the table as soon as we have addressed the epistemological question.

4.6.1 The Epistemological Question

If the inductive premise of the sorites argument is not true, why can we not find any falsifying instances? When faced with our sorites series for "tall," why can we not swipe our fingers through two men in the series and assert that one is "tall" and the other "not tall?" Why does this idea seem so inherently ridiculous?

Let us suppose for the moment that any actual, concrete sorites series physically realizes some canonical sorites series for the context. The following might provide the answer to our questions.

Theorem 6.5 (Law of Quasi-Boundarylessness) Let P be a standard-sensitive predicate, O^P be its densely ordered set of degrees, $\langle K, \langle w, \delta\rangle\rangle$, $K = (W', R', V')$, $W' \subseteq W \times O^P$ be a context, $\langle K', w\rangle \in$ **vagsel** be a constraining selection for $\langle K, \langle w, \delta\rangle\rangle$, and $\langle S, \prec_p^s\rangle$ be a canonical sorites series for P at $\langle K, \langle w, \delta\rangle\rangle$ with a first element α and a last element β. If $\gamma \prec_p^s \eta$, then either $\langle K, \langle w, \delta\rangle\rangle \nVdash DP\eta$ or $\langle K, \langle w, \delta\rangle\rangle \nVdash D\neg P\gamma$.

PROOF: Let $\gamma \prec_p^s \eta$. Now suppose that $\langle K, \langle w, \delta\rangle\rangle \Vdash DP\eta$ and $\langle K, \langle w, \delta\rangle\rangle \Vdash D\neg P\gamma$. By the Law of Quasi-Tolerance (Theorem 5.5), then, $\mu([\gamma, \eta]) \geq f(0)$, and so by definition of $\langle S, \prec_p^s\rangle$, it cannot be that $\gamma \prec_p^s \eta$, contrary to assumption. Thus, either $\langle K, \langle w, \delta\rangle\rangle \nVdash DP\eta$ or $\langle K, \langle w, \delta\rangle\rangle \nVdash D\neg P\gamma$.

This Law of Quasi-Boundarylessness states that for any two degrees directly adjacent in a canonical sorites series for P, one cannot assert of one of the pair that it is P and of the other that it is not P. Now, recall our Law of Precise Boundaries (Principle 4.1). By Definition 6.1, any canonical sorites series for P is a well-ordering. The Law of Precise Boundaries, then, states that our drawing a boundary for P within the canonical sorites series — our finding falsifying instances for the inductive premise — requires there being $\gamma, \eta \in S$ such that $\gamma \prec_p^s \eta$, $\langle K, \langle w, \delta\rangle\rangle \Vdash DP\eta$ and $\langle K, \langle w, \delta\rangle\rangle \Vdash D\neg P\gamma$. Thus, we are unable to draw a boundary anywhere within the canonical sorites series for P; we can find no falsifying instances.

Let us unpack this idea further. Basically, Principle 4.1 and Theorem 6.5 provide an answer to the epistemological question which trades on the nature of the act of finding falsifying instances for the inductive premise. To find such instances requires our being able to assert of two points γ and η adjacent in our sorites ordering that one is P and the other is not. However, it has been a commonplace of dynamic semantics since at least Stalnaker 1978 to hold that the assertion of a proposition within a context requires the context to support the proposition. Given our semantics for the operator D, this simply means that for a speaker to assert within a context $\langle K, \langle w, \delta\rangle\rangle$ that Ψ requires that $\langle K, \langle w, \delta\rangle\rangle \Vdash D\Psi$. Thus, finding falsifying witnesses for the inductive premise requires there being γ and η such that $\gamma \prec_p^s \eta$, $\langle K, \langle w, \delta\rangle\rangle \Vdash DP\eta$, and $\langle K, \langle w, \delta\rangle\rangle \Vdash D\neg P\gamma$, which is just what the Law of Quasi-Tolerance forbids.

In other words, the answer which we offer to the epistemological question is similar to that offered by the account proposed in Graff 2000: one cannot find the "boundary" in the sorites series because, wherever one looks, the principles governing vague language prevent one from locating the boundary there. The way in which a standard-sensitive predicate P depends upon vague selections for its interpretation, combined with the lower bound on the size of vague selections, entails that there is a minimum difference between two objects, below which one can no longer assert of them that one is P and the other not. The construction of a sorites series exploits this fact, and creates a context at which there is a well-ordering in which one cannot draw a boundary for the predicate P. This need not entail, however, that the inductive premise of the sorites argument is true at that context. Why, then, do we think that it's true?

4.6.2 The Psychological Question

Due to the Law of Quasi-Boundarylessness, I cannot imagine a situation in which I could say of two men adjacent in a sorites series for "tall" that one is tall and the other is not tall. Now, usually if I cannot imagine a situation in which I would say that some P is not a Q, common sense dictates to me that "all P's are Q" is true. But what if, in the case of the sorites series, common sense runs afoul of a strange subtlety in our language?

Picture the following scenario: our practices for verifying the truth of sentences in our language are rational simply because they have been discovered in the past to be widely successful. That is, our belief that a universal statement, if false, should admit the discovery of falsifying instances is not based on any *a priori* knowledge of the semantics of universals, but on the fact, discovered at the dawn of prehistory, that in nearly all cases it is beneficial to consider a universal sentence true if no falsifying instances can be found. Perhaps this belief was of such use to our early ancestors that it even found its way into our innate "common sense" concerning the use of language. That belief, like many of our common sense beliefs, need not be true, only extremely useful. What if the sorites argument is, like the Müller-Lyer illusion, a "trick" whereby our deeply-held common sense beliefs are suddenly caught with their pants down.

If the inductive premise of the sorites isn't true, why does common sense compel us to believe it? The answer is that we have adopted a common-sense view of the way in which the semantics of universals interacts with the act of verification, one which the sorites argument interestingly reveals to be incorrect. We are inclined to incorrectly believe the inductive premise because in most everyday cases, in all normal contexts, the inaccuracy of a universal claim entails that it is in principle possible to find objects which one may assert falsify the claim. When one introduces a sorites series, however, one unwittingly manipulates the context to create a singularly bizarre and rare occurrence: a context in which a universal claim may be not be true but it is in principle impossible to find any objects which falsify the claim. This is the real sense in which the sorites paradox witnesses a "breakdown" in the language-game of vague predicates. It is not, however, a case in which the principles governing vague predicates have brought us into inconsistency, but one in which our deep-seated beliefs about how the semantics of universal claims connect with human action become challenged.

This response to the psychological question now permits us to argue that all canonical sorites series should be sorites series by Concept 3.4.

Theorem 6.6 (Law of Sorites Susceptibility) If P is a standard-sensitive context dependent predicate, then P is sorites susceptible.

PROOF: Let P be a standard-sensitive predicate, and let O^P be its densely ordered set of degrees. By Definition 3.3, it is sufficient to show that there is a sorites series for P at some context. So, let $\langle K, \langle w, \delta \rangle \rangle$, $K = (W', R', V')$, $W' \subseteq$

$W \times O^P$ be a context, and $\langle S, \prec_p^s \rangle$ be a canonical sorites series for P at $\langle K, \langle w, \delta \rangle \rangle$ with a first element α and a last element β. The Law of Quasi-Boundarylessness states that for any two degrees $\gamma, \eta \in O^P$ such that $\gamma \prec_p^s \eta$, either $\langle K, \langle w, \delta \rangle \rangle \nVdash DP\eta$ or $\langle K, \langle w, \delta \rangle \rangle \nVdash D\neg P\gamma$. Thus, the speakers at $\langle K, \langle w, \delta \rangle \rangle$ will not be able to assert of any such γ, η that $P\eta$ and $\neg P\gamma$, and so their common sense beliefs regarding the relationship between the truth of universals and the ability to find falsifying instances compel them to believe that at the context $\langle K, \langle w, \delta \rangle \rangle$ the sentence $\forall x, y((Px \wedge y \prec_p^s x) \rightarrow Py)$ is true. That is, to the speakers at the context $\langle K, \langle w, \delta \rangle \rangle$, it appears that $\langle K, \langle w, \delta \rangle \rangle \Vdash D\forall x, y((Px \wedge y \prec_p^s x) \rightarrow Py)$. Concept 3.4, Proposition 6.2, and Proposition 6.3 now imply that $\langle S, \prec_p^s \rangle$ is a sorites series for P at $\langle K, \langle w, \delta \rangle \rangle$.

4.6.3 Interim Summary

Let us pause for the moment and take note of where we stand. Thus far, we have explained our analysis of the sorites paradox, and shown that this analysis can be conjoined with our semantic theory to predict that all standard-sensitive predicates are sorites susceptible.

The resolution which we offer to the sorites paradox is that we believe the inductive premise of the sorites argument to be true because we cannot find any falsifying instances, and we cannot find any falsifying instances because doing so would require our making assertions which our semantic theory says can't happen. Let P be a standard-sensitive predicate. Since there is a lower bound on the sizes of vague selections, there is a minimum difference between two objects such that one may assert of one of the pair that it's P and of the other that it's not. Our perspective on the sorites argument is that it exploits this minimum of difference to create a well-ordering of objects within which one cannot draw a precise boundary for P. At this point, one's common sense beliefs regarding the relationship between the truth of universals and the act of finding falsifying instances kicks in and forces one to the belief that the universal sentence is true, even though it's not. The argument in our proof of Theorem 6.6 shows that this perspective on what is "going on" in the sorites paradox allows our theory to predict that all standard-sensitive predicates are sorites susceptible.

However, our theory has yet to predict that all sorites susceptible predicates are standard-sensitive. Furthermore, we have conspicuously neglected to mention the existence of a difficult variation of the sorites, one which our analysis has yet to face. In the last two sections of this thesis, we address both these matters. We begin, in the next section, with the problem of the "higher-order sorites".

4.7 Quasi-Tolerance and the Higher-Order Sorites

In our discussion of the sorites paradox, we have so far spoken only of the "classical" sorites argument, which operates on a particular vague predicate P from the lexicon of the language. However, recall from our discussion in Section

3.2.2 that for any vague predicate Φ, the predicate "is a borderline case for Φ" is also vague. We may immediately perceive, then, that for any $n \in \mathbb{N}$, a sorites argument may be constructed for the predicate "is an n-th order borderline case for P." The way this realization is often informally put is that, just as there could be no two men with only an inch difference in their height, one of whom was tall and the other not tall, there could be no pair of men with, say, only a half-inch difference in height, one of whom was a borderline case for "tall", and the other not, nor a pair of men with only a quarter-inch difference in height, one of whom was a borderline borderline case for "tall" and the other not,... and so on. Strictly speaking, these are all just versions of the same, old sorites argument, but philosophers sometimes speak of the latter as revealing "higher-order" sorites paradoxes, since they operate on predicates denoting higher-order vagueness classes.

We would like in this section to argue that the general form of analysis offered in the previous section for the classical sorites paradox can likewise be applied to the higher-order sorites paradoxes. That is, we would like not only to offer resolutions to these paradoxes, but also to correctly predict that for all $n \geq 1$, the predicate I^nP is sorites susceptible. Unfortunately, this matter is still largely a topic for future research. Given the simplicity of our formal representations of vague selections, we cannot yet make our arguments regarding the higher-order sorites as precise as those laid out in the last section for the classical paradox. Nevertheless, the following remarks should make clear to the reader how the intended extension is to proceed.

The basic observation informing our approach to the higher-order sorites is that the analysis of Section 4.6 rests only on the presence of a lower bound on the size of the set of $sel(x)$-points in the vague selection. This, when combined with the second Vagueness Principle, establishes the crucial lower bound on the distance between two degrees α, β such that one could in a context be DP and the other $D\neg P$. Similarly, then, one should be able to use the lower bound on the size of vague selections to deduce for any $n \geq 1$ a lower bound on the size of the set of $I^n sel(x)$-points, one which could then establish a minimum distance between two degrees such that one could in a context be DI^nP and the other $D\neg I^nP$. Such a minimum distance would thus permit in any context the construction of a canonical sorites series for the predicate I^nP. Our analysis of the classical sorites would then apply straightforwardly to sorites series for all higher order vagueness predicates I^nP, and we would correctly predict that all these predicates are sorites-susceptible.

Again, although we cannot argue as precisely as we did in previous sections, the following considerations should indicate to the reader how a finer, more sophisticated representation of vague selections might lead to a developed account of the higher-order sorites. Suppose that we have a vague selection of points along a dense linear ordering.

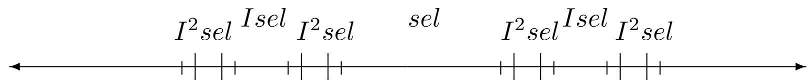

Suppose, moreover, that α and β are two points in the ordering such that for this vague selection $DI^n sel(\alpha)$ and $D\neg I^n sel(\beta)$. Since $DI^n sel(\alpha)$, we may conclude that $I^n sel(\alpha)$ and for no $m \in \mathbb{N}$ is it true that $I^m I^n sel(\alpha)$. Thus α is in some interval of $I^n sel$ points, and it is not in any of the intervals of $I^m I^n sel$ points. We may picture this as follows.

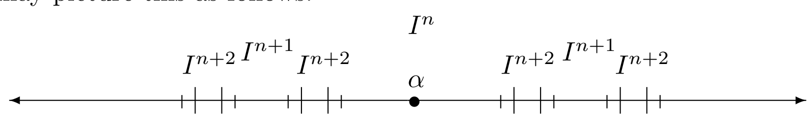

Similarly, from the fact that $D\neg I^n sel(\beta)$, we may conclude that $\neg I^n sel(\beta)$ and for no $m \in \mathbb{N}$ is it true that $I^m \neg I^n sel(\beta)$. Recall that by our semantics for I, $I\Phi$ holds iff $I\neg\Phi$ does. Thus, β is not included in any $I^n sel$ interval, and nor does it lie in any of the intervals of $I^m I^n sel$ points. We conclude, then, that the points α and β must be separated by some interval γ containing for every $m \in \mathbb{N}$ a subinterval of $I^m I^n sel(\alpha)$ points. That is, we know that the following picture depicts the least possible distance between α and β.

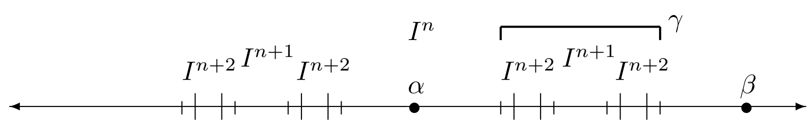

Now, following our discussion in Section 4.5.2, we conclude that each subinterval within γ must have a minimal size. Thus, the interval γ separating α and β must have a minimal size. Let ϵ be the lower bound for $\mu(\gamma)$. Since $[\alpha, \beta] \supset \gamma$, $\mu([\alpha, \beta]) \geq \mu(\gamma) \geq \epsilon$. Therefore, we have deduced an $\epsilon \in \mathbb{R}$ such that $DI^n sel\alpha$ and $D\neg I^n sel\beta$ implies $\mu([\alpha, \beta]) \geq \epsilon$.

The graphic underneath Proposition 5.3 depicts the intuitive relationship within a context $\langle K, \langle w, \delta \rangle \rangle$ between the various vagueness orders for some standard-sensitive predicate P in $\langle K, \langle w, \delta \rangle \rangle$ and the uncertainty orders of a constraining selection for $\langle K, \langle w, \delta \rangle \rangle$. A more sophisticated formalization of vague selections should be able to capture this relationship in full, but according to our depiction in the graphic, all $DI^n sel$ degrees in the vague selection correspond to $DI^{n+1}P$ degrees in the context, and all $D\neg I^n sel$ degrees correspond to $D\neg I^{n+1}P$ degrees. Thus, our lower bound ϵ on the distance between $DI^n sel$ and $D\neg I^n sel$ degrees becomes a lower bound on the distance between $DI^{n+1}P$ and $D\neg I^{n+1}P$ degrees. We conclude, then, that for all $n \geq 1$ there exists an $\epsilon \in \mathbb{R}$ such that $DI^n P\alpha$ and $D\neg I^n P\beta$ implies $\mu([\alpha, \beta]) \geq \epsilon$.

Now, given this value ϵ, it should again be possible to construct a canonical sorites series, this time for the predicate I^nP. Let $\langle K, \langle w, \delta\rangle\rangle$ be any context. By the nature of vague selections and Proposition 5.3, we are guaranteed the existence of points α, β such that $\langle K, \langle w, \delta\rangle\rangle \Vdash DI^nP\alpha$ and $\langle K, \langle w, \delta\rangle\rangle \Vdash D\neg I^nP\beta$. Clearly, there exists some well-ordering $\prec_p^s$ of degrees from O^P such that this α is its first member, β is its last member, and all points γ, η adjacent in the ordering are such that $\mu([\gamma, \eta]) < \epsilon$. By reasoning which should be familiar to the reader, it will thus be impossible for any speakers in $\langle K, \langle w, \delta\rangle\rangle$ to assert of degrees adjacent in this ordering that one is I^nP and the other $\neg I^nP$. If one accepts the analysis of the sorites propounded in the last section, it follows that speakers in $\langle K, \langle w, \delta\rangle\rangle$ should be compelled then to believe that the context supports the untrue inductive premise $\forall x, y((I^nPx \wedge y \prec_p^s x) \rightarrow I^nPy)$, thus creating the existence of a sorites paradox for the predicate I^nP.

Therefore, this envisioned analysis, when fully developed, would correctly predict that for all $n \in \mathbb{N}$ and standard-sensitive predicates P, the predicate I^nP is susceptible to a sorites paradox. Moreover, it would also expand our resolution of the classical sorites to its higher-order cousins.

4.8 Sorites Susceptibility Equivalent to Standards

In this last section, we would like to finally argue that that the analysis of the sorites proposed in Section 4.6, when combined with our theory of predication from Section 3.8.2, predicts that predicates are susceptible to sorites paradoxes only if they are standard-sensitive.

Recall that we hypothesize that for any natural language $\mathcal{L}$, there are two classes of predicates **vague** and **precise**, that every predicate P of $\mathcal{L}$ is a member of exactly one of these classes, and its semantics entirely determined by the Principles of Predication (Principle 3.8.6). We would like now to show that for any $P \in$ **precise**, there cannot be a sorites paradox for P. However, given our present formulation of what it means for there to be a sorites paradox for P, this becomes too easy. Recall that in Definition 3.2, we stated that there is a sorites paradox for P at the context $\langle K, \langle w, \delta\rangle\rangle$ if there exists a sorites series for P at $\langle K, \langle w, \delta\rangle\rangle$, and that this was defined in 3.4 as a certain series of *degrees* of P. Since, by Principle 3.8.6, only vague predicates have an associated degree ordering, there can be no sorites series for any $P \in$ **precise**, and so we obtain the non-sorites-susceptibility of precise predicates for free.

This, of course, is too facile an explanation, and does not truly satisfy what we intend by our desire that the account predict sorites paradoxes not to arise for any $P \in$ **precise**. Again, what we intuitively want is to rule out is the existence of sorites paradoxes for P as they are first informally introduced in Section 1.1 and formulated in Concept 3.1. That is, for any $P \in$ **precise** and any context $\langle K, \langle w, \delta\rangle\rangle$, we would like to rule out the possibility of there being some subset S of the domain at $\langle K, \langle w, \delta\rangle\rangle$, and well-ordering Q of S such that S has a first element α, a last element β, $\langle K, \langle w, \delta\rangle\rangle \Vdash DP\alpha$, $\langle K, \langle w, \delta\rangle\rangle \Vdash D\neg P\beta$, and it appears to the speakers in $\langle K, \langle w, \delta\rangle\rangle$ that $\langle K, \langle w, \delta\rangle\rangle \Vdash D\forall x, y((Px \wedge Qxy) \rightarrow$

Py).

Now, recall that for standard-sensitive predicates, such orderings Q are possible because of the special features of the canonical sorites series. It appears to the speakers in a context that the context supports the sentence $\forall x, y((Px \wedge y \prec_p^s x) \rightarrow Py)$ because one cannot find any two entities γ, η such that $\eta \prec_p^s \gamma$, $\langle K, \langle w, \delta \rangle \rangle \Vdash DP\gamma$, and $\langle K, \langle w, \delta \rangle \rangle \Vdash D\neg P\eta$. To show that precise predicates are not sorites susceptible, then, we want to show that for any context $\langle K, \langle w, \delta \rangle \rangle$ and ordering Q of its domain, there are two entities γ, η such that $Q\gamma\eta$, $\langle K, \langle w, \delta \rangle \rangle \Vdash DP\gamma$ and $\langle K, \langle w, \delta \rangle \rangle \Vdash D\neg P\eta$. Since one can assert in the context that P holds of γ and $\neg P$ holds of η, then by our proposed answer to the psychological question, it should be obvious to the speakers in the context that the universal inductive premise is false. Because speakers *can* draw a boundary for P within the ordering Q, it does not appear to them that $\langle K, \langle w, \delta \rangle \rangle \Vdash D\forall x, y((Px \wedge Qxy) \rightarrow Py)$, and so no sorites paradox ensues. Intuitively, this is correct. No sorites paradox ever arises for a precise predicate because we *know* that the universal inductive premise is false for such predicates. If we could prove such a fact for precise predicates and well-orderings of domains, then we would have a rather intuitive explanation of the non-sorites susceptibility of precise predicates.

This goal, however, is out of the question: it could very well be that in some context $\langle K, \langle w, \delta \rangle \rangle$, we simply know absolutely nothing about the extension of $P \in$ **precise**, and so for no α in the domain of $\langle K, \langle w, \delta \rangle \rangle$ is it the case that $\langle K, \langle w, \delta \rangle \rangle \Vdash DP\alpha$. Similarly, it could be that our factual knowledge in the context $\langle K, \langle w, \delta \rangle \rangle$ simply happens to be such that we don't know for any two objects α, β in the domain that $Q\alpha\beta$, $DP\alpha$, and $D\neg P\beta$.

This latter point is vividly illustrated by the following scenario, which also casts doubt upon the resolution to the sorites offered in Section 4.6. Suppose that you tell me that there is such a thing in topology as a "kablungi space," but you also withhold from me the definition of this concept. Now, given my familiarity with topology, I know that "kablungi space" is a precise concept. Moreover, imagine that you claim to have a set of spaces arranged in a well-ordering Q [13] such that the first space ordered in Q is a kablungi space and the last is not a kablungi space. Let "K" be a predicate denoting kablungi spaces, α be the first space in your ordering, β be the last space, and $\langle K, \langle w, \delta \rangle \rangle$ be the context in which this is all played out. It follows that $\langle K, \langle w, \delta \rangle \rangle \Vdash DK\alpha$ and $\langle K, \langle w, \delta \rangle \rangle \Vdash D\neg K\beta$. Now, since I don't know what a kablungi space is, α and β are the only two objects for which I have definite knowledge whether they are kablungi spaces. Thus, presuming that there are other spaces in Q interpolated between α and β, it follows that for no two spaces γ, η such that $Q\gamma\eta$ do I have $\langle K, \langle w, \delta \rangle \rangle \Vdash DK\gamma$ and $\langle K, \langle w, \delta \rangle \rangle \Vdash D\neg K\eta$.

We conclude that $\langle K, \langle w, \delta \rangle \rangle$ is a counterexample to the conjectured claim that if P is a precise predicate, then for *any* well-ordering Q of the objects of the domain, there are objects α, β such that $Q\alpha\beta$, $\langle K, \langle w, \delta \rangle \rangle \Vdash DP\alpha$ and $\langle K, \langle w, \delta \rangle \rangle \Vdash D\neg P\beta$. Moreover, $\langle K, \langle w, \delta \rangle \rangle$ seems to be a challenge to our

[13] Perhaps you've drawn them all up on a series of cards.

analysis of the sorites proposed in section 4.6. Notice that the situation for "kablungi space" in the context $\langle K, \langle w, \delta \rangle \rangle$ is of the exact same kind as that which supposedly brought about our accepting the inductive sorites premise for the predicate "tall." Our analysis was that a sorites paradox arises for a standard-sensitive predicate P in a context $\langle K, \langle w, \delta \rangle \rangle$ when there was a well-ordering $\langle S, \prec_p^s \rangle$ such that for no two adjacent elements γ, η in the ordering did $\langle K, \langle w, \delta \rangle \rangle \Vdash DP\gamma$ and $\langle K, \langle w, \delta \rangle \rangle \Vdash D\neg P\eta$. Due to our common-sense folk-beliefs regarding how the semantics of universals interfaces with human action, such a context was claimed to "fool" us into accepting the untrue inductive premise. Why, then, do we *still* not accept the inductive premise for "kablungi space" in the imagined context $\langle K, \langle w, \delta \rangle \rangle$?

Clearly, we must slightly amend our analysis of the sorites paradox. Let us now try a somewhat more complicated analysis, one which builds on that proposed in Section 4.6. It will be argued that this analysis reveals the crucial difference between vague predicates such as "tall" and precise predicates like "kablungi space," and correctly predicts that we never accept the inductive sorites premise for precise predicates. Our more complicated analysis, then, will be the one to predict that predicates are sorites susceptible only if they are standard-sensitive.

This "more complicated" analysis, however, is not all that complicated. What I would like to propose is that the key to a predicate's being sorites-susceptible is there existing an ordering of the objects in the domain such that one cannot even *imagine* there being a context in which one could find a boundary for the predicate within the series. That is, the reason I do not in the imagined context $\langle K, \langle w, \delta \rangle \rangle$ accept the inductive premise for "kablungi space" is that I could imagine a context in which I knew more about the meaning of "kablungi space" and so could assert of two objects adjacent in the ordering that one is a kablungi space and the other not. In a sorites series for the predicate "tall", however, I am inclined to accept the inductive premise because there is no imaginable context in which I could assert of two entities adjacent in the ordering that one is tall and the other not. Since I cannot even *imagine* discovering the inductive premise to be false, I am for that reason compelled to believe that it is true.

To make this line of argument precise, let us first establish the following two facts about vague and precise predicates.

Proposition 8.1 Let $\langle K, \langle w, \delta \rangle \rangle$ and $\langle F, \langle v, \chi \rangle \rangle$ be distinct contexts, $P \in$ **vague**, O^P be the densely ordered set of degrees for P, and $\langle S, \prec_p^s \rangle$ be a canonical sorites series for P in $\langle K, \langle w, \delta \rangle \rangle$. For all $\gamma, \eta \in O^P$, if $\eta \prec_p^s \gamma$, then either $\langle F, \langle v, \chi \rangle \rangle \nVdash DP\gamma$ or $\langle F, \langle v, \chi \rangle \rangle \nVdash D\neg P\eta$.

PROOF: Suppose that $\langle F, \langle v, \chi \rangle \rangle \Vdash DP\gamma$ and $\langle F, \langle v, \chi \rangle \rangle \Vdash D\neg P\eta$. By the Law of Quasi-Tolerance (Theorem 5.5), then, $\mu([\eta, \gamma]) \geq f(0)$ and so by the definition of the canonical sorites series for P in $\langle K, \langle w, \delta \rangle \rangle$, it cannot be that $\eta \prec_p^s \gamma$.

Theorem 8.2 (Vagueness is Standard-Sensitivity II) Let $\langle K, \langle w, \delta \rangle \rangle$ be a context, $P \in$ **precise**, and Q be a well ordering of the domain $\mathbb{D}$ at $\langle K, \langle w, \delta \rangle \rangle$ such that Q has a last element β. There exists a context $\langle F, \langle v, \chi \rangle \rangle$ and $\gamma, \eta \in \mathbb{D}$ such that $Q\gamma\eta$, $\langle F, \langle v, \chi \rangle \rangle \Vdash DP\gamma$ and $\langle F, \langle v, \chi \rangle \rangle \Vdash D\neg P\eta$.

PROOF: We will establish the existence of at least one such $\langle F, \langle v, \chi \rangle \rangle$ by first constructing a particular member of **vagsel**. Let S be a subset of $\mathbb{D}$ such that $\beta \notin S$. Let $\langle F', v \rangle$ be an admissible pointed Kripke model with the same domain as $\langle K, \langle w, \delta \rangle \rangle$, such that $F' = (W, R, V)$, where $u, v \in W$ are such that $\mathbf{P}(u) = \mathbf{P}(v) = S$, $R = \{\langle v, v \rangle, \langle v, u \rangle\}$, and $\{\delta' : \langle F', v \rangle \Vdash \mathit{sel}(\delta')\} \neq \{\delta' : \langle F', u \rangle \Vdash \mathit{sel}(\delta')\}$.

Clearly, $\langle F', v \rangle \in$ **vagsel**; the skeptical reader may note that it is simply the model from page 78. Now, let $\langle F, \langle v, \chi \rangle \rangle$, $F = \langle W', R', V' \rangle$, $W' \subseteq W \times O^P$ be a context such that $\langle F', v \rangle$ is a constraining selection for it. Thus, $\langle F', v \rangle$ and $\langle F, \langle v, \chi \rangle \rangle$ have the same domain and so Q is again a well ordering on the domain of $\langle F, \langle v, \chi \rangle \rangle$. Thus, let γ be the last element of S ordered in Q. That is, let $\gamma \in S$ be such that $\exists x(Q\gamma x \vee Qx\gamma)$ and for no $s \in S$, $Q^T\gamma s$. Now, since by assumption $\beta \notin S$ and β is the last element ordered in Q, we know that there exists some $\eta \in \mathbb{D}$ such that $\gamma \neq \eta$ and $Q\gamma\eta$. However, since $\eta \notin S$, then $\eta \notin \mathbf{P}(v)$ and $\eta \notin \mathbf{P}(u)$. Let $\langle t, \psi \rangle \in W'$ be such that $R'\langle v, \chi \rangle \langle t, \psi \rangle$. Since $\langle F', v \rangle$ is a constraining selection for $\langle F, \langle v, \chi \rangle \rangle$, $\langle F, \langle t, \psi \rangle \rangle \nVdash P\eta$. Thus, we derive that $\langle F, \langle v, \chi \rangle \rangle \Vdash D\neg P\eta$. Similarly, since $\gamma \in S$, we have that $\gamma \in \mathbf{P}(v)$ and $\gamma \in \mathbf{P}(u)$, and we conclude that $\langle F, \langle v, \chi \rangle \rangle \Vdash DP\gamma$.

Theorem 8.2 demonstrates that if P is a precise predicate, $\langle K, \langle w, \delta \rangle \rangle$ is a context and Q is a well ordering on its domain, then there exists an imaginable context $\langle F, \langle v, \chi \rangle \rangle$ in which there are two items adjacent in Q, and one can assert of them that P holds of one and $\neg P$ holds of the other. Thus, for any precise predicate and any ordering of the objects in the domain of the context, one can imagine a state of knowledge where they could pick out, identify, assert the presence of a precise boundary in the ordering for the predicate.

This is, however, not the case for standard-sensitive predicates. For any standard-sensitive predicate P, there *is* a well-ordering of the domain such that in *no* imaginable context can one pick out items directly adjacent in the ordering and assert that one is described by P and the other not. If this is quality of an ordering Q which causes speakers to wrongly accede to the inductive premise of the sorites argument, then, again, we predict that all standard-sensitive predicates are sorites susceptible, because in any context we can construct such a Q. Furthermore, we also predict that *no* precise predicates are sorites susceptible, since there can be no such Q for them. Whatever well-ordering one places on the domain, one can easily construct a context at which items γ, η directly adjacent in the ordering are such that one may assert that $P\gamma$ and that $\neg P\eta$.

Now, given our assumption that all predicates in a language may be divided between the classes **vague** and **precise**, we predict that a predicate is sorites susceptible if and only if it is standard-sensitive. Combining this result with our earlier work linking standard-sensitivity to the ability to have borderline cases, we may conclude this thesis with the statement

$$
\begin{gathered}
P \text{ is vague} \Leftrightarrow P \text{ is sorites susceptible} \Leftrightarrow \\
P \text{ has borderline cases} \Leftrightarrow \\
P \text{ is standard-sensitive context dependent}
\end{gathered}
$$

which was to be proved.

Bibliography

[1] Barker, C. 2002. *The Dynamics of Vagueness.* Linguistics and Philosophy, 25. 1 - 36.

[2] Barnes, J. 1982. *Medicine, Experience and Logic.* Barnes, J., J. Brunschwig, M. F. Burnyeat, and M. Schofield (eds.) Science and Speculation. Cambridge University Press, Cambridge.

[3] Black, M. 1937. *Vagueness: An Exercise in Logical Analysis.* Philosophy of Science, 4. 427 - 455.

[4] Bosch, P. 1983. *Vagueness is Context-Dependence.* Ballmer, T. and M. Pinkal (eds.) Approaching Vagueness. North-Holland, Amsterdam.

[5] Cargile, J. 1969. *The Sorites Paradox.* British Journal for the Philosophy of Science, 20. 193 - 202.

[6] Cresswell, M. 1976. *The Semantics of Degree.* Partee, B. (ed.) Montague Grammar. Academic Press, New York.

[7] Dekker, P. 1993. *Transsentential Meditations: Ups and Downs in Dynamic Semantics.* PhD Dissertation. ILLC Dissertation Series 1993-1, University of Amsterdam.

[8] Dummett, M. 1975. *Wang's Paradox.* Synthese, 30. 301 - 324.

[9] Edgington, D. 1992. *Validity, Uncertainty and Vagueness.* Analysis, 52. 193 - 204.

[10] Fine, K. 1975. *Vagueness, Truth and Logic.* Synthese, 30. 265 - 300.

[11] Gaifman, H. 2002. *Vagueness, Tolerance and Logic.* Ms available at http://www.columbia.edu/∼hg17/VTC-latest.pdf

[12] Gerbrandy, J. 1999. *Bisimulations on Planet Kripke.* PhD Dissertation. ILLC Dissertation Series DS-1999-01. Ms available at http://www.illc.uva.nl.

[13] Graff, D. 2000. *Shifting Sands: An Interest-Relative Theory of Vagueness.* Philosophical Topics, 28.

[14] Graff, D. 2001. *Phenomenal Continua and the Sorites.* Mind, Vol. 110 . 440.

[15] Groenendijk, J. and M. Stokhof. 1991. *Dynamic Predicate Logic.* Linguistics and Philosophy, 14. 39 - 100.

[16] Groenendijk, J. and M. Stokhof. 1999. *Meaning in Motion.* von Heusinger, K. and V. Egli (eds.) Reference and Anaphoric Relations. Kluwer, Dordrecht. Ms available at http://turing.science.uva.nl/∼stokhof/papers/mim.pdf

[17] Groenendijk, J., M. Stokhof, and F. Veltman. 1996. *Coreference and Modality.* Lappin, S. (ed.) Handbook of Contemporary Semantic Theory. Blackwell, Oxford.

[18] Hempel, C. 1939. *Vagueness and Logic.* Philosophy of Science, 6. 163 - 180.

[19] Kamp, H. 1975. *Two Theories About Adjectives.* Keenan, E. (ed.) Formal Semantics of Natural Language. Cambridge University Press, Cambridge.

[20] Kamp, H. 1981. *The Paradox of the Heap.* Mönnich, U. (ed.) Aspects of Philosophical Logic. D. Reidel Publishing, Dordrecht.

[21] Kennedy, C. 1997. *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison.* PhD Dissertation. University of California, Santa Cruz. Published (1999) by Garland Press, New York.

[22] Kyburg, A. and M. Morreau. 2000. *Fitting Words: Vague Language in Context.* Linguistics and Philosophy, 23. 577 - 597.

[23] Keefe, R. and P. Smith. 1996a. *Introduction: Theories of Vagueness.* Keefe, R. and P. Smith (eds.) Vagueness: A Reader. MIT Press, Cambridge.

[24] Keefe, R. and P. Smith. 1996b. *Vagueness: A Reader.* MIT Press, Cambridge.

[25] Lappin, S. 2000. *An Intensional Parametric Semantics for Vague Quantifiers.* Linguistics and Philosophy, 23. 599 - 620.

[26] McGee, V. and B. McLaughlin. 1998. *Book Review: "Vagueness" by Timothy Williamson.* Linguistics and Philosophy, 21. 221 - 235.

[27] Mehlberg, H. 1958. *Truth and Vagueness.* From: Mehlberg, H. The Reach of Science. University of Toronto Press, Toronto.

[28] Parikh, R. 1994. *Vagueness and Utility: The Semantics of Common Nouns.* Linguistics and Philosophy, 17. 521 - 535.

[29] Parsons, T. and P. Woodruff. 1995. *Worldly Indeterminacy of Identity.* Proceedings of the Aristotelian Society, 95. 171 - 191.

[30] Peirce, C.S. 1902. *Vague.* J. Baldwin (ed.) Dictionary of Philosophy and Psychology, 748. Macmillan, New York.

[31] Russell, B. 1923. *Vagueness.* Australasian Journal of Philosophy and Psychology, 1. 84 - 92.

[32] Sainsbury, R. M. 1996. *Concepts without Boundaries.* Inaugural Lecture. Reprinted in Keefe and Smith 1996b.

[33] Seuren, P. 1978. *The Structure and Selection of Positive and Negative Gradable Adjectives.* Farkas, D., W. Jacobson, K. Todrys (eds.) Papers from the Parasession on the Lexicon, CLS 14. University of Chicago, Chicago.

[34] Stalnaker, R. 1978. *Assertion.* Cole, P. (ed.) Pragmatics (Syntax and Semantics 9). D. Reidel Publishing, Dordrecht.

[35] Tye, M. 1994. *Sorites Paradoxes and the Semantics of Vagueness.* Tomberlin, J. (ed.) Philosophical Perspectives 8: Logic and Language. Ridgeview Publishing, Atascadero.

[36] Unger, P. 1979. *There Are No Ordinary Things.* Synthese, 41. 117 - 154.

[37] van Benthem, J. 1991. *Language in Action. Categories, Lambdas and Dynamic Logic.* North-Holland, Amsterdam.

[38] van Fraassen, B. 1969. *Presuppositions, Supervaluations and Free Logic.* Lamber, K. (ed.) The Logical Way of Doing Things. Yale University Press, New Haven.

[39] Veltman, F. 1996. *Defaults in Update Semantics.* Journal of Philosophical Logic, 25. 221 - 261.

[40] Veltman, F. 2001. *Het Verschil Tussen* Vaag *en* Niet Precies. Address Upon Confirmation of Professorship. Ms. available at "http:www.uva.nl/onderzoek."

[41] von Stechow, A. 1984. *Comparing Semantic Theories of Comparison.* Journal of Semantics, 3. 1 - 77.

[42] Williamson, T. 1992. *Vagueness and Ignorance.* Proceedings of the Aristotelian Society, Supp. Vol 66. 145 - 162.

[43] Wright, C. 1975. *Language-Mastery and the Sorites Paradox.* Evans, G. and J. McDowell (eds.) Truth and Meaning. Oxford University Press, Oxford.

[44] Wright, C. 1987. *Further Reflections on the Sorites Paradox.* Philosophical Topics, 15. 227 - 290.

If, after I depart this vale, you ever remember me and have thought to please my ghost, forgive some sinner and wink your eye at some homely girl.

H. L. Mencken