THE UNIVERSITY OF CHICAGO

SEMANTIC INDECISION

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE HUMANITIES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF LINGUISTICS

BY
TIMOTHY W. GRINSELL

CHICAGO, ILLINOIS
AUGUST 2017

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

I owe my first debt of gratitude to Chris Kennedy, my advisor and chair. I'm especially grateful for the model of scholarship he set, a model whose primary virtue I'd call generosity. If this dissertation successfully honors the complexity of vagueness phenomena and the richness of the related work, it is because of his example and, inevitably, patience.

I am similarly grateful to members of my committee, Malte Willer and Greg Kobele. They have each improved this dissertation with their insight, rigor, and enthusiasm.

For helpful discussions, advice, and criticism, I'm grateful to Itamar Francez, Saul Levmore, Brian Leiter, Ryan Bochnak, Matt Teichman, Tamara Vardomskaya, Andrea Beltrama, Pete Klecha, Scott AnderBois, Barbara Partee, Angelika Kratzer, Tom Ernst, Emmanuel Chemla, Polly Jacobson, Steve Sloman, Philip Hofmeister, Fabrizio Cariani, Thony Gillies, Anubav Vasudevan, Glen Weyl, Emir Kamenica, Andrew Havens, Michael Morreau, Larry Solum, Andrew Verstein, Omri Ben-Shahar, as well as audiences at SALT 22, the University of Chicago Workshop in Semantics and Philosophy of Language, the Northwestern University Philosophy and Linguistics Workgroup, Brown University's LingLangLunch, and the Public Choice Society's 2015 Annual Meeting.

I am grateful to my family for support and encouragement where they could have been incurious or skeptical, and to the University of Chicago, the Law School, the Jacobs Fund, the Chicago Linguistic Society, and the American Council of Learned Societies for support.

I owe my last debt of gratitude to Michelle.

ABSTRACT

Vagueness is semantic indecision, David Lewis said. This dissertation vindicates this insight by applying social choice theory, the branch of economics concerning collective decision making, to account for linguistic vagueness. Vagueness effects are analogues to "cycles" that sometimes arise in collective decisions. Cycles (e.g. A is preferred to B is preferred to C is preferred to A) may arise in a collective body like a legislature whenever such a body tries to choose among three or more proposals. Though cycles paralyze decision making, the economist Kenneth Arrow proved that (under certain conservative assumptions) cycles are unavoidable.

Cycling is a failure of transitivity, the property that requires if A is related to B, and B is related to C, then A is related to C. Like cycling, vagueness effects depend on intransitivities: if A looks to be as tall as B, and if B looks to be as tall as C, it may not be the case that A looks to be as tall as C. Vagueness effects like the sorites paradox (how many people can move to a little city before it is no longer little?), borderline cases (what is the biggest little city?), and higher-order vagueness (when is a city definitely little, as opposed to just little?) all involve failures of transitivity.

Arrow's result explains these failures. From the view of choice theory, the semantic "choosers" are vague predicates such as healthy, and the "chosen" are different measures by which to determine what counts as healthy. Adjectives like healthy involve evaluations of multiple different criteria in context, such as blood pressure or cholesterol (this is why it is possible to say things like healthy with respect to cholesterol). Speakers compare contextually relevant entities according to these criteria, counting as healthy the entities that rank sufficiently high on sufficiently many, sufficiently important dimensions. These adjectives can therefore be interpreted as choice functions subject to Arrow's result. Replace voters with the different criteria and legislative proposals with contextually relevant entities, and Arrow's result reconstructs itself in the semantics of gradable adjectives.

Vagueness is therefore indecision. And this indecision is semantic in nature: vagueness

effects follow from standard assumptions about the meaning of gradable adjectives plus reasonable assumptions about aggregating many criteria into one. Where many previous approaches to vagueness have focused on the range of the function healthy—the set of truth values—the indecisional approach that I propose focuses instead on this function's domain, most relevantly the set of degrees. When this domain depends on multiple component dimensions (like blood pressure and heart rate for *healthy*), these component dimensions must be aggregated into one "healthiness" scale much in the same way that voters' preferences must be aggregated to produce a collective decision. Vagueness effects follow, and are even expected, from reasonable constraints on this aggregation process.

CHAPTER 1

INTRODUCTION

David Lewis famously argued that "vagueness is semantic indecision," (Lewis, 1986, 213).

> The reason it's vague where the outback begins is not that there's this thing, the outback, with imprecise borders; rather there are many things, with different borders, and nobody has been fool enough to try to enforce a choice of one of them as the official referent of the word "outback."[1]

Semantically, vague predicates like *outback* are choice functions, and they choose among different ways to determine what counts as *outback*. In these types of decisions, it is sometimes impossible to choose—a paradox of choice. That is, at least in some cases, vagueness is a decision problem, and vagueness effects result from paradoxes of choice.

The vagueness effects considered in this dissertation include participation in the sorites paradox (as illustrated in (1)), borderline cases (which is the biggest little city, or the shortest tall man?), and higher-order vagueness (related to the issue of borderline cases: when is someone *definitely tall* as opposed to just *tall*?). These phenomena reflect the difficulty of drawing a boundary between the extension and the anti-extension of a vague predicate—between, for example, the *tall* things and the non-*tall* things.

For example, Reno, Nevada claims to be "the biggest little city in the world." Is this true? Consider the following line of reasoning:

(1) Premise 1. A city with a population of 1 is a little city.

Premise 2. A city with a population of 1 more than the population of a little city is a little city.

Conclusion. A city with a population of 8,000,000 is a little city.

1. Lewis articulated this view in support of the argument that "[t]he only intelligible account of vagueness locates it in our thought and language" rather than in the world.

Iterating the reasoning of the second premise suggests that if Reno is a little city, then so (eventually) is New York.

The reasoning in (1) has gone wrong in at least two ways. First, the indifference relationship implied by the second premise is intransitive. The second premise reflects that vague predicates are "tolerant." Kamp (1981) explains tolerance like this: suppose the objects a and b are indistinguishable in the respects relevant to some property P; then either both a and b satisfy P or neither of them does (2).

(2) Premise 2: $\text{little}(a) \wedge a \sim_{\text{little}} b \rightarrow \text{little}(b)$

A vague predicate like *little* is tolerant in this sense. However, the truth of both the first premise and the conclusion tells us that the relation "indifferent with respect to size", represented by $\sim$, is not transitive. Reno may be observationally indistinguishable from Cincinnati with respect to population, and Cincinnati from Honolulu, but Reno and Honolulu may noticeably differ in size.

A second and related way in which (1) goes wrong is that small changes along one dimension of size (like population) do not translate smoothly into changes in the meaning of the predicate *little*. Small changes in population do not appear to change the value of $[\![\text{little(Reno)}]\!]$ from true to false. Rather, the value of $[\![\text{little(Reno)}]\!]$ appears to change only with sufficiently large changes in population. The mapping from one dimension of *little* to the meaning of *little* appears to be discontinuous.

The sorites paradox therefore demonstrates that vagueness is a threat both to valid reasoning and to the idea that we have a firm grasp on what the terms of our language mean. Vagueness threatens valid reasoning because it encourages us to proceed from apparently true premises to a false conclusion, a procession normally forbidden by the laws of logic. It threatens our grasp on meaning by challenging the presumption of truth-functionality. As van Rooij (2011, 125) explains, "[A] minimal requirement for any theory of meaning seems to be that one knows the meaning of a declarative sentence if one knows under which circumstances it is, or would be, true." But if we do not know which things are *little* and which are not, we are unable to satisfy this "minimal requirement" for a sentence like *Reno is the biggest little city*.

The challenges posed by vagueness are challenges of collective decision making. Both properties demonstrated by the sorites paradox—intransitivity and discontinuity—are problems of collective choice. For example, the Marquis de Condorcet observed that intransitivities (or cycles, e.g. A is preferred to B is preferred to C is preferred to A) may arise in a collective body like a legislature whenever such a body tries to choose among three or more alternatives. In 1950, Kenneth Arrow proved that (under certain conservative assumptions) a there is no collective decision-making procedure that is guaranteed to avoid cycling. Generalizations of Condorcet's work have also shown that collective decisions tend to be discontinuous: if a voter changes her preferences in a minor way, the output of the collective decision may change in a major way.

Vagueness effects are distinguished by intransitivity and discontinuity, and these properties arise in lexical predicates as a result of limitations on collective choice. Evidence for the social choice view comes from properties of "multidimensional" adjectives like *healthy*. These adjectives involve evaluations of multiple different criteria in context, such as blood pressure or cholesterol in the case of *healthy* (this is why it is possible to say things like *healthy with respect to cholesterol*). Speakers compare contextually relevant entities according to these criteria, counting as *healthy* the entities that rank sufficiently high on sufficiently many, sufficiently important dimensions.

These adjectives are therefore collective choice functions, subject to limitations on collective choice. Replace Arrow's voters with the different criteria and Arrow's candidates with contextually relevant entities, and Arrow's impossibility result reconstructs itself in the semantics of gradable adjectives like *healthy*.

Thus, the sorites paradox and the cycling problem are related. The cycling problem is a failure of transitivity: if the legislature prefers proposal A to B, B to C, and C to A, the legislature's preferences are intransitive. Similarly, the sorites paradox is widely though to implicate the intransitivity of the "indifference" relation. It is not true, for example, that if A and B are ranked indifferently with respect to health, and if B and C are ranked indifferently with respect to health, then A and C are ranked indifferently with respect to health. In fact, the intransitivity represented by the sorites paradox and the intransitivity represented by cycling arise from the same source, the aggregation of many judgments into one.

Similarly, a generalization of Arrow's results shows that collective decisions will also be discontinuous under certain circumstances. That is, small changes in the voters' preferences do not necessarily result in a small change in the outcome. For example, Gaertner discusses the following example: consider voters 1-5 with preferences among the alternatives v, x, y, z in descending order, as in (3).

(3)

voter 1	voter 2	voter 3	voter 4	voter 5
x	y	z	x	z
y	v	v	v	x
z	x	y	z	v
v	z	x	y	y

Alternative x beats all the other alternatives by majority vote in a pairwise contest. However, if the order of x and z are reversed in voter 2's preference ranking—voter 2's least prefered alternatives—then z beats all alternatives. A small change in the voters' preferences results in a large change in the outcome to the election. The discontinuity represented by the sorites paradox and the discontinuities intrinsic to collective decision making also arise from the same source: the aggregation of many judgments into one. And other vagueness effects, like borderline cases and higher-order vagueness, likewise flow from the demands imposed by any reasonable collective decision-making function.

Vagueness is therefore indecision. And this indecision is semantic in nature: vagueness effects follow from standard assumptions about the meaning of gradable adjectives plus reasonable assumptions about aggregating many scales into one. There is no answer to the question "is Reno the biggest little city?" because this answer involves a decision made impossible by the semantics of the gradable adjective *little* itself. This is independent of whether *Reno is the biggest little city* is true on some "precisifications" but not others (supervaluationism); or only partly true (many-valued logic); or true but unknowable (epistemicism); or true in reasonably related contexts (contextualism). The indecisional approach to vagueness is compatible with each of these approaches, but it depends on none of them.

* * *

This dissertation is structured as follows. In chapter 2, I use the framework of choice theory to compare four different approaches to vagueness rooted in the philosophical tradition—supervaluationism, epistemicism, many-valued logic, and contextualism—with their linguistic correlates. This comparison reveals a tension between philosophical and linguistic desiderata of a theory of vagueness. Philosophical theories generally abandon a well-behaved domain of degrees in order to invalidate the second premise of the sorites paradox. In contrast, linguistic theories seek to maintain this well-behaved structure, primarily out of a concern for modeling the semantics of comparatives. The choice-theoretic approach reveals that the precise locus of this tension is the relationship between indifference and comparison as defined over the set of degrees.

Chapter 3 shifts the focus of the inquiry into the nature of vagueness from truth to degrees. In this chapter, I first introduce the limitations of collective choice in more detail, and I then construct a parallel between collective decisions and the use of multidimensional adjectives like *healthy*. Since collective decision making is governed by Arrow's Theorem, I argue that multidimensional adjectives are subject to Arrow's Theorem as well. As a consequence of Arrow's result, multidimensional adjectives may be intransitive in their indifference relation, and this intransitivity gives rise to vagueness effects.

Chapter 4 extends this reasoning to unidimensional adjectives. These adjectives are also multicriterial in the sense that adjectival standards—the threshold for determining whether an entity is *tall* or not—depend on multiple types of information. Furthermore, the indecisional approach accounts for the distinction between "relative" and "absolute" adjectives (respectively, adjectives that display vagueness effects and adjectives that do not, like *tall* vs. *flat*). In the social choice approach, absolute adjectives are adjectives whose criteria minimally "agree" on the rankings of the contextually relevant entities. Therefore, they are not subject to collective-choice-induced intransitivities, and they do not display vagueness effects.

Chapter 5 returns to the original analogy between legislative decision making and vagueness phenomena. Just as the workings of a legislature inform the indecisional approach to vagueness phenomena, the indecisional approach informs how we should understand legislatures. In debates over legal interpretation, the cycling problem has standardly been used to argue that legislatures have no intent: since it is impossible to coherently aggregate the intent of individual legislators (for reasons delineated by Arrow), there would seem to be no such thing as legislative intent. The indecisional approach to vagueness shows that this argument cannot be correct. If legislative intent is incoherent because of the cycling problem, then multidimensional adjectives like *healthy* must be similarly incoherent. Since adjectives like *healthy* are obviously meaningful, the attack against legislative intent—an attack relied upon by the United States Supreme Court when it makes decisions—does not succeed.

CHAPTER 2

THEORIES OF VAGUENESS AS THEORIES OF CHOICE

Against the challenges posed by vagueness, philosophers and linguists have mounted a two-pronged defense. The philosophers have attempted to limit the damage to classical logic; the linguists have tried to reconcile vague predicates with a compositional account of other linguistic phenomena, like comparatives. But these efforts have sometimes been in tension with one another. To see how, consider the sorites paradox as translated into logical form (4), assuming $P(x_1)$ is true and $P(x_n)$ is false (van Rooij, 2011).

(4) Premise 1. $P(x_1)$

Premise 2. $\forall x_i, x_{i+1} : (P(x_i) \wedge x_i \sim^P x_{i+1}) \rightarrow P(x_{i+1})$

Conclusion. $P(x_n)$

Philosophers have tended to tinker with Premise 2 (the so-called *tolerance principle*) in order to avoid the conclusion $P(x_n)$. This usually involves blocking the transitivity of the "indifference" relation $\sim^P$, such that $x_1 \sim^P x_2$ and $x_2 \sim^P x_3$ does not imply $x_1 \sim^P x_3$. For linguists, however, transitivity of the comparative relation $\succ^P$ defined over P is essential: if $x \succ^P y$ and $y \succ^P z$, then $x \succ^P z$.

The trouble arises with combinations of the relations, as in (5).

(5) $a \succ^P b$ and $b \sim^P c$

As Sen (1970) shows, if (5) implies that $a \succ^P c$, then the indifference relation is transitive—a result that causes problems for analyses of the sorites paradox. On the other hand, if (5) does *not* imply that $a \succ^P c$, then (5) implies that either $c \succ^P a$ or $a \sim^P c$. If the former, then the comparative relation is intransitive. And if the latter, then it is possible for two entities to differ with respect to property P but to remain unrelated by the comparative relation defined over P, $\succ^P$, a result at odds with demands imposed by the semantics of natural language comparatives.[1]

1. Except, importantly, the semantics of multidimensional adjectives.

In the present chapter, I argue that every theory of vagueness confronts the tension between the vague "positive" use and the precise comparative use. This manifests as the tension between completeness (the idea that every entity is either indifferent to, better than, or worse than every other) and transitivity. Choice theory makes it possible to compare different theories of vagueness according to whether and how they navigate this tension.

An examination of four prominent theories of vagueness—supervaluationism, epistemicism, many-valued logic, and contextualism, and their respective linguistic implementations—then reveals a tendency towards a "choice rule" with a transitive strict relation $\succ$ and an intransitive indifference relation $\sim$. This option is not on Sen's menu above. As explained in chapter 3, such a choice rule only makes sense in the context of a *collective* decision.

2.1 Vagueness phenomena

There is agreement as to the types of core linguistic phenomena a theory of vagueness must capture: participation in the sorites paradox, borderline cases, and higher-order vagueness (Williamson, 1994; Kennedy, 2007). In one way or the other, these phenomena reflect the difficulty of drawing a boundary between the extension and the anti-extension of a vague predicate; between, for example, the *healthy* things and the non-*healthy* things.

The sorites paradox was illustrated above in (1).[2] As the structure of the sorites paradox reveals, it is applicable to a wide range of natural language expressions. Common nouns like *heap* are susceptible, as are gradable adjectives (so-called because they can appear in comparative constructions) like *tall* and *red*, as are adverbs (*very*), quantifiers (*many*), verbs (*start*), proper names (*Chicago*), and definite descriptions (*the border between Illinois and Iowa*).[3]

2. Notably, in order for the second premise of the sorites paradox to work, there must be some scale, or *dimension*, along which one increments (or decrements) values. For adjectives like *tasty* or *familiar*, it is not always clear what that dimension should be. The problems of dimensions are addressed in the next chapter.

3. Kennedy (2007) and others also class context-dependence with the types of phenomena that vague

Vague predicates also have "borderline cases." When a vague predicate is applied to a borderline case, we are not sure if the resulting proposition is true or false. For example, even if we know Clarence's exact height of 1.85m (about 6 feet), the argument goes, we may still be unsure whether Clarence qualifies as *tall* (and hence unsure whether the proposition *Clarence is tall* is true).

Notably, the presence of a borderline case is not sufficient to ensure vagueness. For example, it may be possible to define a predicate *definitely-tall** such that people above 180cm in height fall into its extension, people 170cm or below do not, and we remain non-committal about the rest (Williamson, 1994). The people between 170cm and 180cm in height might fairly be called borderline cases, yet *definitely-tall** is not vague in the same way as its natural-language counterpart *definitely tall*. In particular, while a predicate like *definitely-tall** has an identifiable boundary between the definitely-tall* entities and the borderline cases—180cm—the natural-language *definitely tall* does not. This is the essence of higher-order vagueness: boundarylessness in a predicate's extension.

Thus, vague predicates may be distinguished from non-vague predicates on the basis of how entities are sorted into their denotations. With non-vague predicates, like the mathematical predicate *prime*, it is clear (or could be made clear) for any entity whether it falls within or without the boundaries of the predicate's denotation. In contrast, vague predicates seem to have blurry denotation boundaries. We can identify clear cases of an entity being *tall*, for instance, but we are hard-pressed to identify the shortest tall entity (or the biggest little city).

This sorting problem distinguishes vagueness from ambiguity or underspecification. An ambiguous word has two (or more) semantically distinct meanings, but this does not characterize vague expressions. Indeed, as Lakoff (1970) noted, ambiguous expressions behave differently from vague expressions under verb phrase ellipsis. For example, though the word

predicates display. As I argue in later chapters, context dependence is a property of vague predicates primarily because the dimensions relevant to their interpretation vary with context.

bank is ambiguous between "shore of a river" and "financial institution," the word must be accorded the same meaning in both the antecedent predicate and the elided predicate (6a). With vague expressions, this is not true. In (6b), the meaning of *tall* may be different in each conjunct. For example, *tall* may mean something like "over six feet" for Sam and something like "over four feet" for Tony if Sam is twenty years-old and Tony is three.

(6) a. Stephen went to the bank, and Ruth did too.

b. Sam is tall, and Tony is too.

Nor is vagueness something like generality or underspecification. The utterance *A woman wrote Middlemarch* is perhaps underspecified for the author role, but it is not vague in terms of the sorites or borderline cases (at least not because the proposition uses the phrase *a woman* rather than *George Eliot*) (Williamson, 1994).

2.2 Choice theory

The problem of sorting the healthy entities from the non-healthy entities is equivalent to the problem of choosing the healthy entities from the rest. In this section, I motivate a view of gradable adjectives as "choice functions." This view is similar to that presented by van Rooij (2011, 140). He defines a *context structure* $\langle I, C, V \rangle$ with I a nonempty set of individuals, C the set of finite subsets of I, and a valuation V that assigns to each $c \in C$ those individuals in c that count as being P, for some predicate P. Then, $P(c) = \{x \in c \mid x \in V(P, c)\}$. The idea, as van Rooij explains, is that "P is thought of as a choice function, selecting the *best* elements of c."

The choice-functional approach and the traditional approach to predicate meaning are closely related; in particular, both approaches associate a predicate's meaning with an extension framed in terms of sets. On the traditional approach, the meaning of the word *book* is the set of books, the meaning of the word *lawyer* is the set of lawyers, and the meaning of the word *healthy* is the set of healthy things (Russell, 1905). The assumption that entities can be sorted like this is problematic for vague predicates. After all, *tall* relates to a kind of scale (presumably of healthiness), and it is possible to order entities along a scale, and any finite set that can be ordered has an upper bound and a lower bound.[4] But asking "what is the lower bound for *healthy*?" is like asking "who is the sickest healthy person?"

In addressing the problems presented by vagueness, the choice-functional approach has three advantages over the traditional semantic approach. First, by taking choice behavior as primitive, this approach is explicit about the "choice rule", how a set of entities (the domain of individuals) is sorted into P and its complement. Second, the choice-functional approach makes available a wide array of formal results from rational choice theory and social choice theory, branches of economics concerned with modeling choice behavior. And third, this approach enables the comparison of previous theories of vagueness according to their associated choice rules. This last task is the project of section 2.3.

2.2.1 Rational choice theory

Rational choice theory addresses the modeling of decision behavior in which one agent chooses a preferred option among a range of options according to a consistent criterion.[5] It makes qualitative predictions that tend to be confirmed, it makes these predictions based on a simple model of the chooser and her constraints, and it is widely applicable.

Despite these advantages, rational choice theory's empirical failings in economics and psychology have led to interest in other approaches. Approaches like "bounded rationality," for instance, tend to stress that the decision-maker's choice rests on a much wider array of factors than the chooser's objectives and her constraints—factors like the order in which options are presented, preexisting emotional states, etc.

In this subsection, I will focus on rational choice theory despite its apparent limitations

4. See, e.g., Sen (1970). As many have argued, including Sassoon (2013b), the same problems exist for *lawyer* and *book*.

5. Much of this subsection summarizes Mas-Colell et al. (1995, Chapter 1), especially the formal results.

for two reasons. First, rational choice theory is simple and its, results far-reaching. Second, results from rational choice theory are easily translatable into the the typed lambda calculus frequently used in natural language semantics. It is therefore straightforward to incorporate the insights of rational choice theory into a formal semantic framework.

Let X be a set of available options. In choice theory, this is usually available goods, like bottles of beer and wine, but it could just as easily be something more complex, like proposals for improving civic transportation. Define an agent's weak preferences over these alternatives with the notation R, where P is the strict preference relation and I is the indifference relation.[6]

(7) $aRb \iff aPb$ or aIb

Thus, a is strictly preferred to b iff aPb, and the agent is indifferent between a and b iff aIb.

Preference relations are binary relations,[7] and it will be helpful to catalogue the relevant properties of binary relations for future use. The first two properties, transitivity and completeness, are the most important.

Definition 2.2.1. (**Transitivity**) A binary relation R is transitive iff $xRy \wedge yRz \rightarrow xRz$ for all x, y, z.

Definition 2.2.2. (**Completeness**) A binary relation R is complete iff $xRy \vee yRx$ for all x, y.

Definition 2.2.3. (**Reflexivity**) A binary relation R is reflexive iff xRx for all x.

Definition 2.2.4. (**Asymmetry**) A binary relation R is asymmetric iff $xRy \rightarrow \neg yRx$ for all x, y.

6. I try to stick to the symbols in Roberts (1979) for consistency. However, since I use P in other contexts, I will usually represent the strict preference relation by $\succ$. In addition, since other authors have used $\sim$ to represent indifference, I will often use it as well.

7. That is, a subset of the Cartesian product $A \times A$, for some set A.

Definition 2.2.5. (**Antisymmetry**) A binary relation R is antisymmetric iff $xRy \wedge yRx \to x = y$ for all x, y.

The weak preference relation R is commonly assumed to be reflexive and transitive, and a binary relation satisfying these properties is called a *pre-order*. It is possible to define other binary relations based on some combination of the properties above.

Definition 2.2.6. (**Pre-Order**) A binary relation R is a pre-order iff it is **reflexive** and **transitive**.

(8) The relation of set inclusion $\subseteq$ is reflexive and transitive, but not complete.

Definition 2.2.7. (**Weak Order**) A binary relation R constitutes a weak order on some domain iff it is **transitive** and **complete**.

(9) The relation "at least as tall as" constitutes a weak order on the set of people.

Definition 2.2.8. (**Total Order**) A binary relation R is a total order iff it is **transitive**, **complete**, and **antisymmetric**.

(10) The relation $\geq$ is a total order on the set of real numbers.

Thus, the relation $\subseteq$ is a pre-order but not a weak order because it is not complete. The relation "at least as tall as" is a weak order but not a total order because it is not antisymmetric.[8] And so forth.

2.2.2 The preference relation

Within rational choice theory, it is possible to define an agent's choice behavior in terms of a preference relation and, conversely, to define a preference relation in terms of choice behavior. Given some modest assumptions, these approaches yield equivalent results.

8. *John is at least as tall as Clarence* and *Clarence is at least as tall as John* does not imply that John is Clarence.

For the first approach, let the set of possible, mutually exclusive alternatives be X. Define a binary preference relation $\succsim$ on the set of alternatives X, where $\succsim$ allows comparison between the members of X. $x \succsim y$ is read as "x is at least as good as y" ($\succsim$ is therefore a weak preference relation, see 2.2.1). Define the strict and indifferent preference relations as follows:

(11) a. Strict preference: $x \succ y$ iff $x \succsim y$ but not $y \succsim x$.

b. Indifference: $x \sim y$ iff $x \succsim y$ and $y \succsim x$.

A preference relation is considered to be *rational* just in case it is complete and transitive.

Definition 2.2.9. (**Rational preference**) A binary relation R is rational iff it is **transitive** and **complete**.

In other words, a rational preference relation is a weak order. As Mas-Colell et al. (1995, 7) explain, transitivity "goes to the heart of the concept of rationality." Transitivity implies that it is impossible for an agent's preferences to *cycle*: for the agent to prefer a to b, b to c, and then c to a.

Preferences can fail to be transitive, as with "just noticeable differences". I provide an example taken from Mas-Colell et al. (1995).

> If we ask an individual to choose between two very similar shades of gray for painting her room, she may be unable to tell the differences between the two colors and may therefore be indifferent. If we offer the individual the choice between the lighter of the two shades and a slightly lighter shade of gray, she may again be unable to tell the difference. We can continue in this fashion moving to progressively lighter shades, and the agent will express indifference at each step. Yet if we offer her the choice between the darkest shade in our sequence and the lightest, she will be able to distinguish between them and would likely prefer one to the other.

Vague predicates display this type of effect.

Other counter-examples to transitivity involve cognitive biases like framing effects. Framing effects turn on information to which the rational choice model is insensitive. The following famous experiment by Kahneman & Tversky (1984) is an example (as described by Mas-Colell et al. (1995)).

> Tverksy and Kahneman asked experimental subjects to answer three questions.
>
> (12) Suppose you are about to buy a stereo for \$125 and a calculator for \$15. You learn there is a \$5 calculator discount at another store branch, ten minutes away. Do you make the trip?
>
> (13) You learn there is a \$5 stereo discount at another store branch, ten minutes away. Do you make the trip?
>
> (14) You learn both items are out of stock. You must go to the other branch, but as compensation you will get a \$5 discount. Do you care which item is discounted?
>
> The fraction of people answering "yes" to the first question is much higher than the fraction answering "yes" to the second. Furthermore, if we expect indifference in response to the third question, the set of answers is intransitive. The option "traveling to save \$5 on a calculator" is preferred to the option "buying both items at the current store." The option "buying both items at the current store" is preferred to the option "traveling to save \$5 on a stereo." But the last question suggests that the options "traveling to save \$5 on a calculator" and "traveling to save \$5 on a stereo" are indifferent with respect to one another. This is profile of preferences violates transitivity (a is preferred to b, b is preferred to c, but a is indifferent with respect to c).

From a complete and transitive rational preference relation $\succsim$, it is possible to construct a choice function by maximizing according to that preference relation (15).[9]

(15) For any nonempty $B \subseteq X$, $C(B, \succsim) = \{x \in B \mid x \succsim y \text{ for all } y \in B\}$

The elements in the set defined by (15) are the decision maker's most preferred alternatives.

2.2.3 Choice behavior

There is another way to go about modeling choice that takes choice behavior, rather than a preference relation, as primitive. Define a *choice structure* as the following (Mas-Colell et al., 1995):

(16) $(\mathcal{B}, C(\cdot))$, where

a. $\mathcal{B}$ is a family of nonempty subsets of X; every element of $\mathcal{B}$ is a set $B \subset X$, and

b. $C(\cdot)$ is a choice rule that assigns a nonempty set of chosen elements for every set B; that is, $C(B) \subset B$.

For example, a choice structure is defined in (17).

(17) $(\mathcal{B}, C_1(\cdot))$, where $X = \{x, y, z\}$ and $\mathcal{B} = \{\{x, y\}, \{x, y, z\}\}$. The choice rule $C_1(\cdot)$ may be defined as $C_1(\{x, y\}) = \{x\}$ and $C_1(\{x, y, z\}) = \{x, y\}$.

If we are only interested in choice behavior that is consistent or reasonable in some sense, then we can impose restrictions on this behavior. For example, we would be surprised if an individual chose x from the set $\{x, y\}$ but y from the set $\{x, y, z\}$. The Weak Axiom of Revealed Preference ("WARP") is one way to avoid this type of anomalous choice behavior. A choice structure satisfies WARP if (18) holds (Mas-Colell et al., 1995).

(18) If for some $B \in \mathcal{B}$, with $x, y \in B$, we have $x \in C(B)$, then for any B' such that $x, y \in B'$, and $y \in C(B')$, then we must also have $x \in C(B')$.

9. If X is finite, this choice set is guaranteed to be non-empty. I will ignore the infinitary case for now.

In prose, if x is ever chosen when y is available, then there can be no set containing both elements from which y is chosen but x is not. WARP also implies the following definition of a *revealed preference relation*:

Definition 2.2.10. (**Revealed preference relation**) Given the choice structure $(\mathcal{B}, C(\cdot))$, a revealed preference relation $\succsim^*$ is defined as $x \succsim^* y$ iff there is some $B \in \mathcal{B}$ such that $x, y \in B$ and $x \in C(B)$.

The Weak Axiom of Revealed Preference and the conditions imposed on rational preferences (completeness and transitivity) result in very similar relations $\succsim$ and $\succsim^*$. Indeed, it turns out that a rational preference relation necessarily generates a choice structure that obeys WARP. Moreover, if a choice structure satisfies WARP *and* if $\mathcal{B}$ includes all subsets of X, then there is a rational preference relation consistent with the choice structure's selections.

2.2.4 Predicates as choice functions

These equivalences point out some consequences of treating predicates as choice functions, as van Rooij (2011) has. For example, let the denotation of *happy* in some context c be as in (19).

(19) $[\![\text{happy}]\!]^c = \{x \mid \forall y \in c,\ xRy\}$

In order to solve the riddle of vagueness, we want to know what properties the relation R has—what is the choice rule for happiness?

The results from the previous subsection suggest that if the choice rule R is interpreted as the rational preference relation $\succsim$ (a weak order), then R will not model vagueness effects. This is because such a relation is transitive in its indifference relation and will validate the sorites paradox. For example, suppose $[\![\text{happy}]\!]^{\{x,y\}} = \{x, y\}$ and $[\![\text{happy}]\!]^{\{y,z\}} = \{y, z\}$. Then $x \sim y$, and $y \sim z$. By the transitivity of $\succsim$ (and hence $\sim$), $x \sim z$. Since the sorites paradox is invalid, a rational preference relation fails to respect the semantics of vague predicates.

It is also possible to start with a strict preference relation $\succ$ and to derive indifference and the relation $\succsim$ from there (20).

(20) a. Indifference: $x \sim y$ iff $\neg x \succ y$ and $\neg y \succ x$

b. Weak preference: $x \succsim y$ iff $\neg y \succ x$

Importantly, this rule gives rise to an intransitive indifference relations in some cases. In choosing between bundles of goods (like beer and wine), an individual may prefer 10 cans of beer and 2 bottles of wine to 9 cans of beer and 2 bottles of wine: $(10, 2) \succ (9, 2)$. However, such an individual may be indifferent as between these two options and 5 cans of beer and 3 bottles of wine: $(10, 2) \sim (5, 3)$ and $(9, 2) \sim (5, 3)$. Note, then, that both $(9, 2)$ and $(10, 2)$ are indifferent to $(5, 3)$, but $(10, 2) \succ (9, 2)$. One of the latter two bundles is strictly preferred to the other, but both are indifferent to a third.

This entails the intransitivity of the indifference relation, and such intransitivity is plausible when the ranked entities have multiple attributes that figure into the ranking (like wine and beer). As discussed below, the choice rule adopted by theories of vagueness includes a transitive strict preference relation and an intransitive indifference relation along the lines of (20). Such a choice rule suggests that the set of available options X is composed of bundles—that is, the choice rule is "multidimensional." Chapter 3 elaborates on this in detail.

2.3 Theories of vagueness as theories of choice

This section examines four prominent responses to the challenge of vagueness–supervaluationism, epistemicism, many-valued logic, and contextualism[10]–and their linguistic implementations.

10. This list leaves out what Williamson (1994) calls "nihilism": the idea that vague expressions are not meaningful and that natural language is inconsistent. I omit this response for two reasons. First, this

It falls short of a survey of the literature on vagueness.[11] Yet a clear trend emerges: these prominent theories have negotiated (with varying degrees of success) the tension between an intransitive indifference relation $\sim$ and a transitive comparative relation $\succ$.

2.3.1 Supervaluationism

Supervaluationism, like other responses that abandon classical logic, has sought some way to invalidate the tolerance principle, repeated in (21).

(21) $\forall x_i, x_{i+1} : (P(x_i) \wedge x_i \sim^P x_{i+1}) \rightarrow P(x_{i+1})$

The supervaluationist approach rejects the validity of the tolerance principle by acknowledging truth-value gaps (Fine, 1975; Kamp, 1975). Accordingly, some uses of the predicate P are true (falling into the positive extension of *tall*, $[\![\text{tall}]\!]^+$, for example), some are false (falling into $[\![\text{tall}]\!]^-$), and some are neither. The main advantage of this approach is that it can accommodate vague predicates while also preserving the valid formulas of classical logic.

Supervaluationism represents the intuition that we can assign a predicate's borderline cases to either its extension or its anti-extension, as long as we are consistent. For example, suppose that $i = 1, \ldots, 100$ grains of sand definitely do not make a heap. Suppose that $j = 201, \ldots, 300$ grains of sand definitely make a heap. Within the remaining values $n =$

response leads to massive incoherence. For example, sorites paradoxes may be constructed for both of the predicates *rich* and *poor*. If the conclusion of each sorites argument is valid (i.e. a person with \$0 is rich and a person with \$1 billion is poor), then these predicates either apply to everyone or no one (if the first person is rich, then the last person is rich; if the first person is poor, then the last person is poor). If the predicates apply equally to everyone, then this attempt to accommodate vagueness leads to inconsistency (*John is both rich and poor*). If the predicates apply to no one, then expressions like *rich* and *poor* do not increase the referential power of the language to which they belong. Then why have them? This type of theorist cannot answer "because they are useful," because the term *useful* is subject to its own sorites series, and therefore it, too, is true of nothing. Second, there has been no work in the nihilist tradition on the interaction of vagueness with the comparative construction and other linguistic phenomena.

11. See, e.g., Keefe (2000).

$101 \ldots 200$, it is possible to draw the line between heaps and non-heaps at any n. These middle values are the borderline cases. The meaning of *heap* is reflected in the class of all *acceptable* values of n; that is, the meaning of a vague predicate is just all the acceptable ways of making it precise.

A sentence is "super-true" if it is true on all acceptable fully precise interpretations. (The super-true sentences include the valid formulas of classical logic.) A sentence is "super-false" if it is false on all acceptable fully precise interpretations. To see how this solves the paradox of the heap in our toy example: the sentence *i grains make a heap* is super-false, the sentence *j grains make a heap* is super-true, and the sentence *n grains make a heap* is true on some acceptable interpretations and false on others. The tolerance principle is therefore super-false: on each acceptable interpretation, there is some n such that n grains do not make a heap but $n+1$ grains do.

Kamp (1975) applies supervaluationism to a compositional theory of gradable adjectives. Kamp begins his treatment with a response to Venneman & Bartsch (1972), who propose a norm-based account of vague adjectives. On their account, *tall* takes its meaning from the comparative relation *taller than*. The adjective *tall* denotes the set of entities that are taller than most of the others in some domain.

(22) $[\![\text{tall lawyer}]\!] = \{i \in [\![\text{lawyer}]\!] \mid \text{for most } i' \in [\![\text{lawyer}]\!], i \text{ is taller than } i'\}$

Kamp noted several problems with such an approach. First, a denotation like (22) excludes cases in which most (or even all) lawyers are, in fact, tall. Second, a meaning like (22) pinpoints the cutoff betwen the positive and negative extensions of adjectives like *tall* with implausible precision, as Kennedy (2007) notes. Indeed, it is still possible to construct a sorites sequence even when the heights of all the lawyers in the domain of individuals are known. However, if the semantics of *tall* were norm-based (that is, dependent on the average height, or the height of a majority of entities, etc.), then such a sorites series should not be possible. Anything above the norm should count as *tall*, while anything below the norm should count as *not tall*. Third, Kamp objected to what he calls the "primacy" of the comparative form over the positive form (*tall*) that (22) seems to suggest.

Kamp's answer to the challenge posed by vague adjectives involves rejecting the comparative-first notion. Instead, he takes adjectives like *tall* to represent partial information: "the semantic interpretation is relative to information states (contexts) in which predicate denotations are only partially specified" (Sassoon, 2013b, 66). The denotation of a predicate like *tall* is separated into a set of individuals falling (in a context) under the positive denotation, the set of individuals falling under the negative denotation, and those falling under neither denotation.

Kamp's idea is that, as information is added to a context the denotation "gap" shrinks, and more entities get sorted into the positive or negative extension. But the gap has to shrink in an orderly way. For example, suppose the "ground" context c consists of entities in the positive extension of *tall*, $[\![\text{tall}]\!]^{+}$, entities in the negative extension of *tall*, $[\![\text{tall}]\!]^{-}$, and some entities in neither. Each fully specified context is compatible with a number of ways of fully classifying the entities in the extension gap; we can call these extensions *total extensions*. In a total extension, each entity is sorted into either the positive or negative denotations. However, if two entities differ in height, then the if the shorter of the two is put into the positive extension, the taller of the two is put into the positive extension as well. This sort of constraint is well represented by Klein's Consistency Postulate (Klein, 1980).[12]

(23) For entities a, b and predicate P, if a counts as P in [context] c, and if $b \geq a$ with respect to the ordering of the domain of P, then b counts as P in c. Similarly, for a [context] c' in which a counts as $\neg P$ and if $a \geq b$, then b counts as $\neg P$ in c'.

The total extensions represent "different ways to fix the cutoff points between the tall and non-tall entities." (Sassoon, 2013b, 69). This cutoff is usually called the *standard*. For example, in (24), each t_i represents a different standard–a cutoff point–between tall and not tall entities (like 2 meters, 1.95 meters, etc.).

12. I have replaced Klein's use of comparison classes with "contexts" in order to maintain consistency with Kamp's presentation.

(24) 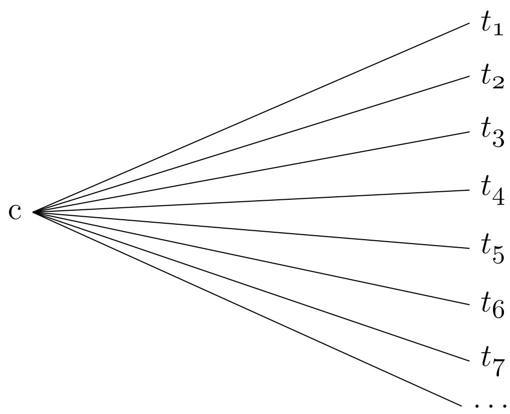

(Sassoon, 2013b, 69)

Thus, if Ruth is on the borderline of being tall, *Ruth is tall* will be true on some total extensions and false on others. On the other hand, if *Ruth is tall* is true in the ground context c, then (due to the constraints on how information is added to the context), *Ruth is tall* is super-true. This solves the sorites paradox, because the second premise is super-false; there is no one cutoff point between the *tall* entities and the non-*tall* entities on all total extensions.

Kamp then defined comparatives in terms of the positive form (i.e. *taller* in terms of *tall*). Under Kamp's theory, the comparative *Sonia is taller than Elena* is true if the set of total extensions in which *Elena is tall* is true is a proper subset of the set of total extensions in which *Sonia is tall* is true (see 25). That is, Sonia meets more standards of tallness than Elena does; this will only be the case when Sonia's height is greater than Elena's height.[13]

(25) $[\![\text{Sonia is taller than Elena}]\!] = 1 \text{ iff } \{t \mid [\![\text{Elena is tall}]\!]^t = 1\} \subset \{t' \mid [\![\text{Sonia is tall}]\!]^{t'} = 1\}$

Supervaluationism's choice rule

The consistency postulates put forward by supervaluationist theories (26a) bear a strong resemblance to WARP (26b).

13. Kamp's paper is important for many other reasons. One reason of particular linguistic significance is his argument for treating vague adjectives as intersective rather than subsective predicates, in contrast to Montague's (1974) approach.

(26) a. "[S]uppose that of two persons u_1 and $u_2 \ldots u_1$ is [more intelligent] than $u_2 \ldots$ There is no model M' such that u_2 falls into [the extension of *intelligent* and u_1 does not]." (Kamp, 1975, 137)

b. A choice structure satisfies WARP if for some $B \in \mathcal{B}$, with $x, y \in B$, we have $x \in C(B)$, then for any B' such that $x, y \in B'$, and $y \in B'$, then we must also have $x \in C(B')$.

Intuitively and formally, these postulates amount to the same thing. We know that if a choice structure satisfies WARP, there is a binary relation $\succsim$ that rationalizes that choice structure. But then the relation $\succsim$ is a complete and transitive ordering over the domain B. Kamp's postulate (26a) is an application of WARP to the case where B represents the extension of some predicate.

Klein's (1980) consistency postulate, repeated below, also imposes a weak order on the entities in the context, as he himself showed.

(27) For entities a, b and predicate P, if a counts as P in [context] c, and if $b \geq a$ with respect to the ordering of the domain of P, then b counts as P in c. Similarly, for a [context] c' in which a counts as $\neg P$ and if $a \geq b$, then b counts as $\neg P$ in c'.

The consistency postulate is therefore equivalent to imposing WARP.

While (26a) and (27) seem to *presuppose* a weak order through the relations *taller than* and $\geq$, respectively (Kennedy, 2001), it is possible to state the restrictions in both (26a) and (27) with reference only to denotation membership, as van Benthem (1982) shows in (28). I have stated van Benthem's axioms in terms of "comparison classes," context-dependent local domains according to which the standard is determined.

(28) For all models M, all individuals a, b in the domain of individuals D, and comparison class $X \subseteq D$ such that $[\![\mathrm{P(a)}]\!]^{M,X} = 1$ and $[\![\mathrm{P(b)}]\!]^{M,X} = 0$,

a. **No Reversal**: There is no $X' \subseteq D$ such that $[\![\mathrm{P(b)}]\!]^{M,X'} = 1$ and $[\![\mathrm{P(a)}]\!]^{M,X'} = 0$;

b. **Upward Difference**: For all X', if $X \subseteq X'$, then there is some $c, d \in [\![\text{P}]\!]^{M,X'}$ such that $[\![\text{P(c)}]\!]^{M,X'} = 1$ and $[\![\text{P(d)}]\!]^{M,X'} = 0$;

c. **Downward Difference**: For all X', if $X' \subseteq X$ and $a, b \in X'$, then there is some $c, d \in X'$ such that $[\![\text{P(c)}]\!]^{M,X'} = 1$ and $[\![\text{P(d)}]\!]^{M,X'} = 0$.

In words, No Reversal states that if a is P in some comparison class and b is not, there is no comparison class in which b is P and a is not. As Burnett (2014, 15) puts it, Upward Difference holds that if there is a P/not-P contrast in a comparison class X, there there is such a contrast in every larger comparison class. Finally, Downward Difference states that if there is a contrast between a and b with respect to P in some comparison class, this contrast is preserved in every smaller comparison class that contains both a and b.[14] It is possible to prove that these axioms also imply a weak order for the comparative, e.g., *taller than.* But if a binary relation (like a comparative) imposes a weak order, it is equivalent to the order imposed by WARP.[15]

Therefore, the supervaluationist choice rule rejects the transitivity of the indifference relation; on each total extension, there is some cutoff x_{i+1} such that $x_i \sim^P x_{i+1}$, but $P(x_i)$ and $\neg P(x_{i+1})$. For any sufficiently rich total extension, then, there will be some triple x_{i-1}, x_i, x_{i+1} such that $x_{i-1} \sim^P x_i$ and $x_i \sim^P x_{i+1}$ but not $x_{i-1} \sim^P x_{i+1}$ (instead, the comparative relation will hold between x_{i-1} and x_{i+1}). This approach also preserves the transitivity of the comparative relation, thanks to consistency postulates like (27).

However, the supervaluationist approach suffers from a number of other problems. First, the supervaluationist gives up bivalence, the idea that every statement in the object language is either true or false. While bivalence survives on each fully precise interpretation, this seems to imply (as Dummett (1975) points out) that if we only had more information, we could

14. See Kennedy (2011) for problems that van Benthem's (1982) axioms have with respect to *implicit comparison*, discussed in 2.3.2.

15. Assuming, again, that every entity is related to every other by either the comparative relation $\succ$ or the indifference relation $\sim$.

pin down the sharp cutoff between a vague predicate's extension and its anti-extension. But it is unlikely that more research will reveal the precise cutoff between the tall and not-tall entities.

Second, the supervaluationist gives up compositionality, the idea that the meaning of an expression is determined by the meaning of its constituent parts and how they are put together. Let $heap(n)$ stand for the proposition that n grains make a heap. The conjunction $(heap(n) \vee \neg heap(n))$ is super-true because it is true in all total contexts. However, neither of the disjuncts is super-true. Similarly, the existential proposition *there exists an n such that n grains make a heap but $n+1$ grains do not* is super-true, but it lacks a super-true instance. If truth is supertruth, then supervaluationism sacrifices compositionality.

Similarly, if validity is the preservation of truth, then this lack of compositionality also poses a danger to validity. Let "definitely p" be true if and only if p is super-true (following Williamson (1994)). The inference from p to "definitely p" is valid because it preserves super-truth: if p is super-true, then "definitely p" is super-true. However, the formula "if p then definitely p" is not valid: if p is true but not super-true, then the conditional has a true antecedent and a false consequent. Put another way, the deduction theorem for classical logic is not valid in supervaluationist logic. In supervaluationism, truth-preserving compositions of formulae may not be super-truth-preserving, which in turn infects the ability to reach conditional conclusions.

Third, as Kennedy (2001); Klein (1980) and others have noted (and as Kamp himself intimated), the supervaluationist approach rules out *a priori* contexts in which all the entities in a context do not count as, e.g., *tall*. If all the entities in the ground context c are already sorted into $[\![\text{tall}]\!]^-$, there is no total extension of this context that includes anything but non-tall entities. Relatedly, the gap-based account does not discriminate between two entities that are definitely tall. If both entities are tall in the ground context c, then they will continue to be tall in the total extensions of c. But then the definition of comparatives in (25) is broken, because the set of total extensions in which $[\![a \text{ is tall}]\!]$ is true is equal to the set of total extensions in which $[\![b \text{ is tall}]\!]$ is true, though a and b may be of different heights.

Fourth, in light of the choice-functional analysis, Kamp's concern about whether the comparative relation $\succ$ or the predicative choice function $P(\cdot)$ is "primitive" seems misguided. As the results from rational choice theory demonstrate—at least for complete and transitive binary relations—it is possible to generate an equivalence between both approaches.

These problems have led linguists to focus instead on theories that associate adjectives with *degrees*. Degree-based approaches depend on a different response to vagueness, epistemicism.

2.3.2 Epistemicism

Suppose we are in a total extension of a supervaluationist context, but we don't know *which* total extension. Then there is a sharp cutoff between the positive and negative extensions of vague terms; we just don't know what it is.[16] This is the situation posited by epistemicism (Williamson, 1994; Sorensen, 2001).

As its title suggests, epistemicism is not primarily a semantic theory. Rather, epistemicism traffics in what we can possibly know about the meanings of the terms we use. Accordingly, Williamson and others draw from epistemic principles to assess the puzzle posed by vagueness. One of these principles is the *margin for error* principle (29).

(29) For "S knows that p" to be true (in a situation s), p must be true in any marginally different situation s' (where one forms the same beliefs using the same methods) in which "S believes p" is true. (Williams, 2012)

The margin for error principle is intended to enforce the idea that only *reliable* true beliefs count as knowledge. If p is false at some very similar situation s', then S's true belief that p in s is too lucky to count as knowledge.

Williamson (1994, 217) uses the example of an observer looking out over a crowded

16. This helpful segue is from Sassoon (2013b).

stadium. Suppose that there are exactly 20,000 people in the stadium, and suppose the observer has the true belief that there are not exactly 10,000 people in the stadium. In any marginally different situation–say where there are 19,999 people or 20,001 people in the stadium–the observer's belief will remain true. In Williamson's terms, the observer's belief is reliable. However, if (in the original situation) the observer has the true belief that there are not 19,999 people in the stadium, the observer's belief will *not* remain true in marginally different situations.[17] The observer's belief is "too risky to constitute knowledge" (Williamson, 1994, 226).

If meaning supervenes on use, as Williamson accepts, then the margin for error principle works to dissolve the sorites paradox. Suppose slightly different patterns of use may determine slightly different cutoffs for a predicate like *bald*. Then a true belief like *Sam is bald* does not count as knowledge when Sam is a borderline case of baldness: a slightly different pattern of use (which we could not tell apart from the current pattern of use) could have placed the cutoff for baldness such that Sam fell into the anti-extension of *bald*.

For the epistemicist, then, the sorites paradox fails because the second premise is false: there is a sharp cutoff between the *bald* and the *not bald* entities. Due to the margin for error principle, however, we are not in a position to know where that cutoff is. Epistemicism is therefore able to preserve principles like bivalence and compositionality along with the other advantages of an approach to natural language semantics based on classical logic.

Perhaps because of these advantages, many modern linguistic treatments of gradable adjectives have adopted the epistemic approach (Barker, 2002; Kennedy, 2007). Degree-based semantics is the most prominent of such treatments. The degree-based approach takes comparison as its starting point. Comparisons involve points on a scale, and a scale is an abstract representation of measurement, usually taken to comprise a dense linearly ordered set of points, or "degrees." The ordering of degrees is relativized to a dimension–some

17. It is false when the stadium contains exactly 19,999 people.

gradable property like *height*, *density*, *health*, etc.[18]

"Once scales and degrees are introduced into the ontology, it becomes possible to analyze gradable predicates . . . as expressions that relate objects in their domains to degrees on a scale" (Kennedy, 1999, 65). On this approach, gradable predicates are of a different type than non-gradable predicates (Cresswell, 1976a; von Stechow, 1984; Heim, 1985). In denoting a relationship between entities and degrees, a gradable predicate like *tall* has a denotation like (30).[19]

(30) $[\![\text{tall}]\!] = \lambda d \lambda x.\text{height}(x) \geq d$

In this case $[\![\text{tall}]\!] = 1$ if the entity x meets or exceeds a contextually determined standard of height (represented in (30) by d). A degree phrase like *five feet* is said to saturate the degree variable d such that *Elena is five feet tall* if Elena's height meets or exceeds the degree denoted by *five feet*. By treating the degree variable as an argument, these analyses are able to account for a wide variety of degree morphology, like *five feet tall*, *very tall*, and much besides.[20]

Kennedy's (2007) approach catalogues the advantages of adopting an epistemicist view of vagueness. His explanation for vagueness effects is based on the "scale structure" of gradable adjectives. First, Kennedy observes that the difference in vagueness effects between so-called "relative" gradable adjectives like *expensive* and absolute gradable adjectives like *open* correlates with the difference between "open scale" and "closed scale" gradable adjectives. These terms derive from earlier work on the scale structure of gradable adjectives (Kennedy

18. The notion of a dimension will play a large role in later chapters. I note here that some properties taken as scalar primitives, like *health*, may actually be combinations of more basic properties.

19. Within the camp of degree-based approaches to vagueness, there is a significant division between the "relational" analysis, which treats gradable predicates as relations between entities and degrees as in (30), and the "measure theoretic" analysis, which treats gradable predicates as measure functions that must combine with a morpheme POS to form the positive. My comments here are intended to apply to both.

20. In particular, this analysis recognizes quantification over degrees, a concept important in many theories of comparison (Heim, 1985; Kennedy, 1999; von Stechow, 1984; Schwarzschild, 2008).

& McNally, 2005; Rotstein & Winter, 2004). This work shows that some gradable adjectives are compatible with modifiers like *perfectly*, while other gradable adjectives are compatible with modifiers like *slightly*, as shown in (31).

(31) a. ??perfectly/??slightly {tall, deep, expensive, likely}

b. ??perfectly/slightly {bent, bumpy, dirty, worried}

c. perfectly/??slightly {straight, flat, clean, unworried}

d. perfectly/slightly {full, open, opaque}

By adopting a semantics of degree for gradable adjectives, Kennedy & McNally were able to make sense of this data. On the degree-based account, the degrees associated with gradable adjectives are formalized by a triple $\langle D, <, \delta \rangle$ including the set of degrees D, an ordering on this domain $<$, and a dimension δ that provides the property to be measured (for instance, cost in the case of *expensive*). Kennedy & McNally explained the data in (31) by arguing that gradable adjectives may differ as to the structure of D, and in particular whether D does or does not have maximal or minimal elements.

For example, if an adjective's associated set D includes a maximal element, the adjective should be compatible with modifiers that pick out maximal degrees on a scale, like *perfectly* (see 31c). Similarly, if an adjective's associated set D includes a minimal element, the adjective should be compatible with modifiers that allude to minimal degrees on a scale, like *slightly* (see 31b). And if the set D includes neither maximal nor minimal elements, the adjective should be compatible with neither modifier (as in 31a), while if the set D contains both maximal and minimal elements, the adjective should be compatible with both modifiers (as in 31d). The following schematizes this "typology of scale structures":

(32) *A typology of scale structures* (Kennedy, 2007, 33)

OPEN	(tall, expensive)	○———○
LOWER CLOSED	(bent, dirty)	●———○
UPPER CLOSED	(straight, flat)	○———●
(TOTALLY) CLOSED	(full, open)	●———●

Kennedy & McNally also observed a correlation between scale structure and vagueness effects. If the D associated with the adjective contains at least one endpoint, the adjective tends to lack vagueness effects. Otherwise, it displays them. Kennedy (2007) further explicated this correlation: closed-scale adjectives typically have interpretations associated with their endpoints. For instance, where two objects differ in length (though both may be short or long when judged individually), experiment subjects readily accepted a definite description with a relative adjective like *long*, as in (33).

(33) Pass me the long one.

In contrast, subjects rejected this use with closed-scale adjectives like *full*. If neither of two cups was filled to the top, even though the cups contained differing amounts of liquid, subjects rejected sentences like (34) (Syrett et al., 2010).

(34) Pass me the full one.

These results indicate that closed-scale adjectives like *full* have endpoint-oriented interpretations. This also explains why closed-scale adjectives tend to lack vagueness effects. There is a sharp cutoff between endpoint degrees and the non-endpoint degrees, and it is easy to tell on what side of the cutoff a particular entity falls. Borderline cases and sorites arguments are therefore less acceptable with adjectives that have endpoint-oriented interpretations. In (35), for example, the second premise is not as readily acceptable as it is in (1).

(35) Premise 1: A rod that has 10 degrees of bend is bent.

Premise 2: A rod that is 1 degree less bent than a bent rod is bent.

Conclusion: A rod that has 0 degrees of bend is bent.

These observations led Kennedy to propose an explanation of vagueness as a function of “standing out.” What it means to “stand out” is to be on the upper end of a natural transition, a transition from a minimal degree to a non-minimal degree (as in the case of *bent*) or from a non-maximal degree to a maximal degree (as in the case of *full*) (Kennedy, 2007, 32). Since relative adjectives lack minimum and maximum degrees, their associated notion of “standing out” must derive from context and is therefore inherently unstable.

Kennedy represented this relation in the gradable adjective’s “standard of comparison”: an entity is *tall*, for instance, if its degree of height exceeds the degree necessary to stand out in context.

(36) $[\![\text{tall}]\!] = \lambda x.\text{tall}(x) \geq \mathbf{s}(\text{tall})$

a. The function **s** chooses a standard of comparison in such a way as to ensure that the objects that *tall* is true of “stand out” in the context of utterance, relative to the kind of measurement that the adjective encodes.

b. “x is tall” is true iff x stands out in context relative to the comparison class.

If Kennedy’s definition of standing out is right, then closed-scale adjectives (those with endpoints to their scale structure) always have a potential degree that stands out. Open-scale adjectives, on the other hand, do not. Kennedy then adopted a principle of “Interpretive Economy” that says that if a stand-out degree is part of the scale structure, that degree is usually chosen by the degree function.[21] This correctly derives the facts. Closed-scale adjectives usually do not display vagueness effects because their “stand-out” relation is fixed by Interpretive Economy.

Kennedy’s approach also accounts for the phenomenon of implicit comparison (what Kennedy calls the *crisp judgment* effect). Imagine two books, one 500 pages long and one 499 pages long. In such a scenario, the explicit comparative in (37a) is felicitous, but the implicit comparative in (37b) is infelicitous.

21. The actual formulation of Interpretive Economy is less ad hoc: “Maximize the contribution of the conventional meanings of the elements of a sentence to the computation of its truth conditions” (Kennedy, 2007, 36).

(37) a. This book is longer than that book.

b. # This book is long compared to that book.

The difference between the two judgments comes down to the stand out relation: since lengths of 500 pages and 499 pages do not stand out with respect to one another, the use of the relative adjective *long* to distinguish between them is infelicitous.[22]

Epistemicism's choice rule

The choice rule for an epistemicist account is unknowable, but we do know this: it draws a boundary between a vague predicate's positive and negative extensions. This boundary-drawing is the basis on which epistemicists falsify the second premise of the sorites paradox.

Tellingly, however, the epistemicist's margin for error principle–the basis for the stand out relation–is compatible with an intransitive indifference relation $\sim$, at least on one interpretation. Luce (1956) demonstrated the usefulness of the margin for error for measurement theory. Luce's stated goal was to capture the intransitivity of the notion of indifference or indistinguishability. As Luce writes,

> A person may be indifferent between 100 and 101 grains of sugar in his coffee, indifferent between 101 and 102, . . . , and indifferent between 4999 and 5000. If indifference were transitive he would be indifferent between 100 and 5000 grains, and this is probably false.

Luce captured this idea in the form of a *semi-order*, whose axioms appear in (38).

(38) a. $\forall x : \neg R(x, x)$

b. $\forall x, y, v, w : (R(x, y) \wedge R(v, w)) \rightarrow (R(x, w) \vee R(v, y))$

c. $\forall x, y, z, v : ((Rx, y) \wedge R(y, z)) \rightarrow (R(x, v) \vee R(v, z))$

22. See Kennedy (2011); Van Rooij (2011) for more arguments over whether supervaluationism can account for crisp judgments.

Put in terms of measurement theory, in a semi-order, $x \succ y$ is true iff there is some real-valued function f and some fixed margin of error ϵ such that $f(x) > f(y) + \epsilon$. Two entities may be indifferent with respect to some property if they both fall within the margin of error ϵ. Epistemicism's choice rule may therefore be intransitive, at least in the indifference relation.

On the semi-order interpretation, epistemicism's choice rule R is also transitive in its comparative relation. The problem is that the semi-order interpretation implies that two entities may differ with respect to P and yet fail to be related by the comparative $\succ^{P}$. The comparative relation $\succ$ is more nearly something like "significantly more than," but such a relation is likely unacceptable as an interpretation of the comparative in English. For example, both of the sentences in (39) may be true at once, yet the semi-order as defined by Luce does not allow for this.

(39) a. Sam is tall, and if Sonia is 1mm shorter than Sam, Sonia is tall.

b. Sam's height is 1mm greater than Sonia's; therefore, Sam is taller than Sonia.

Perhaps because of this problem, linguistic implementations of epistemicism have instead assumed that while the the positive form includes a margin of error, the comparative lacks such a margin. Kennedy, for example, accomplishes this by associating the positive form with a silent morpheme POS that contributes the margin for error.

Notably, though, the dense, totally ordered scales presupposed by degree theories (see Kennedy (1999); von Stechow (1984); Heim (2001); Kamp & Partee (1995) among others) are incompatible with the semi-order interpretation. Such scales are total orders, not semi-orders. The binary relation R defined over this scale is transitive and complete, which means that the indifference relation is transitive and complete as well. If an intransitive indifference relation is invoked by the use of the positive form—and examples like (39) suggest it is—then the scales undergirding degree semantics are too rich to directly represent semi-ordered phenomena. A semi-order allows indifference ($\sim$) where a total order requires comparison ($>$).

Furthermore, it is not clear that such information-rich scales are necessary to represent adjectival meanings. Most adjectival predicates seem to only encode weak-order-type rankings. For example, we do not associate fine-grained ordering information with degrees of happiness or beauty (though, admittedly, we do associate these sorts of orderings with height and weight and other extensive measures) (Sassoon, 2013b, 87).[23] Nor is it clear how to order entities with respect to health when the various measures of health vary widely. For example, if Nino has high blood pressure and low cholesterol, and Stephen has low blood pressure and high cholesterol, should we place them on a par?

Three more questions confront the epistemicist account. First, this account fails to explain why we are disposed to treat the sorites reasoning as valid (Fara, 2002). If we know there is a boundary between heaps and non-heaps, we should dismiss the sorites premise out of hand. But we don't.[24] Second, linguistic implementations of epistemicism are inextricably tied up with the semantics of gradable adjectives. These implementations do not seem to extend to other lexical categories that may display vagueness effects, like common nouns (of which *heap* is one) (van Rooij, 2011). Finally, and most importantly, the epistemicist has no account of how meaning supervenes on use in such a way as to draw sharp boundaries between the extensions and anti-extensions of vague predicates. As Williamson (1994, 209) explains, meaning may supervene on use in an "unsurveyably chaotic" way. An inquiry that leads to unsurveyable chaos is an inquiry at its end.

23. The information associated with each view—ordinal rankings in weak orders and cardinal rankings in degree theories—is related to the existence of vagueness phenomena (3). This explanation favors the weak-order view.

24. Smith (2008, 36) offers one explanation: "Our ignorance of where the cut-off is makes us think there is no cut-off at all. That is why we are inclined to accept the inductive premiss, even though, according to the epistemicist, the inductive premiss is actually false."

2.3.3 Many-valued logic

Unlike epistemicism, many-valued logic tinkers with the machinery of classical logic. It does this by adding truth values, resulting in systems ranging from three truth values to an infinite number of truth values (Hyde, 2008). This subsection discusses the theory put forward by Cobreros et al. (2012) and implemented linguistically by Burnett (2014). This system is equivalent to a three-valued logic (Cobreros et al., 2015).

Cobreros et al.'s approach begins with the classical notion of satisfaction in first order logic. The satisfaction of a well-formed formula is defined over structures of interpretation $\langle D, I\rangle$ that include a domain of entities (D) and an interpretation function (I).

(40) a. For a constant a, $I(a) \in D$.

b. For a predicate P, $I(P) \subseteq D$.

Cobreros et al.'s innovation was to add to this structure a binary indifference relation $\sim$ for every predicate P. They call this structure $\langle D, I, \sim\rangle$ a *T-model* (for *tolerant*). The indifference relation $\sim$ is symmetric ($x \sim y \rightarrow y \sim x$) and reflexive ($x \sim x$) but possibly non-transitive.

Unlike supervaluationist and epistemicist approaches, Cobreros et al. set out to validate the tolerance principle, repeated in (41), which features prominently in the sorites paradox.

(41) $\forall x \forall y[P(x) \wedge x \sim y \rightarrow P(y)]$

"If some individual x is P, and x and y are only imperceptibly different in respects relevant for the application of the predicate P, then y is P as well" (Cobreros et al., 2012, 348).

In order to validate the tolerance principle without also validating the sorites paradox, Cobreros et al. must alter classical logic.

In the authors' system, the definition of classical truth remains the same as ever (42), where the notation $\phi[d/x]$ stands for the result of substituting d for every free occurrence of x in ϕ.

(42) *c-truth*

$M \vDash {}^c P(a)$ iff $I(P)(a) = 1$

$M \vDash {}^c \neg\phi$ iff $M \nvDash {}^c \phi$

$M \vDash {}^c \phi \wedge \psi$ iff $M \vDash {}^c \phi$ and $M \vDash {}^c \psi$

$M \vDash {}^c \forall x \phi$ iff for every d in D, $M \vDash {}^c \phi[d/x]$

Classical validity is therefore preserved as *c-validity*, even in T-models.

In order to validate the tolerance principle, Cobreros et al. introduce *t-truth* (tolerant truth) and its dual *s-truth* (strict truth).[25]

(43) a. *t-truth*

$M \vDash {}^t P(a)$ iff $\exists d \sim a : M \vDash {}^{\mathbf{c}} P(d)$

$M \vDash {}^t \neg\phi$ iff $M \nvDash {}^{\mathbf{s}} \phi$

$M \vDash {}^t \phi \wedge \psi$ iff $M \vDash {}^t \phi$ and $M \vDash {}^t \psi$

$M \vDash {}^t \forall x \phi$ iff for every d in D, $M \vDash {}^t \phi[d/x]$

b. *s-truth*

$M \vDash {}^s P(a)$ iff $\forall d \sim a : M \vDash {}^{\mathbf{c}} P(d)$

$M \vDash {}^s \neg\phi$ iff $M \nvDash {}^{\mathbf{t}} \phi$

$M \vDash {}^s \phi \wedge \psi$ iff $M \vDash {}^s \phi$ and $M \vDash {}^s \psi$

$M \vDash {}^s \forall x \phi$ iff for every d in D, $M \vDash {}^s \phi[d/x]$

Note that a sentence $P(a)$ is t-true if there is *some* entity d related by the indifference relation $\sim$ to a such that $P(d)$ is c-true. Similarly, a sentence $P(a)$ is s-true if *every* entity d related by the indifference relation $\sim$ to a is such that $P(d)$ is c-true. As the authors note, the indifference relation $\sim$ is "rigid," and therefore "borderline cases of a predicate are definite on the present approach" Cobreros et al. (2012, 354).

How does this validate the tolerance principle? The tolerance principle is t-valid (44) (that is, for every T-model M, $M \vDash {}^t \phi$).

25. As the authors note, introducing just *t-truth* produces the wrong results.

(44) $\vDash {}^t\forall x\forall y\forall z[P(x) \wedge x \sim y \wedge y \sim z \rightarrow P(z)]$

To prove this conditional, it suffices to show that the antecedent is false in all T-models or the consequent is true. If the antecedent is false, then $\nvDash {}^t[P(a) \wedge a \sim b \wedge b \sim c]$, which, by the duality of tolerant and strict truth, means $\vDash {}^s[P(a) \wedge a \sim b \wedge b \sim c]$. If $\vDash {}^sP(a)$, then $\forall d \sim a$, $\vDash {}^cP(d)$. So in particular $\vDash {}^cP(b)$. Since $\vDash {}^tP(c)$ iff $\exists d \sim b$ such that $\vDash {}^cP(d)$, and since $\vDash {}^cP(b)$ and $b \sim c$, we have $\vDash {}^tP(c)$. (The case in which the consequent is true is trivial.) This approach also t-validates apparent contradictions, which Cobreros et al. contend are felicitous for borderline cases.

(45) a. Ruth is both tall and not tall.

b. The theatre is both empty and not empty.

Burnett (2014) extends Cobreros et al.'s system to account for the distinction between absolute and relative adjectives. Burnett's account of the relation between absolute and relative gradable adjectives hinges on properties of the indifference relation $\sim$.[26] On her account, the indifference relation is symmetric for relative adjectives ($a \sim b$ implies $b \sim a$) but not for absolute adjectives (46).

(46) Axioms obeyed by the indifference relation $\sim$ for relative and absolute gradable adjectives:

a. **Symmetry**: For all relative adjectives P, $a \sim {}^Pb$ implies $b \sim {}^Pa$

b. **Asymmetry**: For all absolute adjectives Q and entities a, b, if $[\![Q(a)]\!] = 1$ and $[\![Q(b)]\!] = 0$, then $b \nsim a$

The intuition behind this is the following: imagine two towels, one completely dry (a) and one with a few drops of water on it (b). Burnett claims that we are more likely to assent

26. Burnett divides the class of absolute adjectives into total absolute adjectives and partial absolute adjectives, but this complication is irrelevant to the point I'm making here. For expository convenience, I will present only the analysis for what she calls total absolute adjectives, which I will call simply "absolute adjectives."

to the proposition that the slightly wet towel resembles the completely dry towel in terms of dryness than to the proposition that the completely dry towel resembles the slightly wet towel in terms of dryness–that is, $b \sim a$ but not $a \sim b$.[27]

Combining (46) with a few other axioms, Burnett is able to explain the differences in context sensitivity (and, on her account, therefore vagueness) between relative and absolute gradable adjectives. As the example above suggests, absolute predicates may sometimes display the types of borderline cases and sorites paradoxes associated with vagueness effects, as in (47).

(47) Premise 1. A (large) theatre with 2 people in it is empty.

Premise 2. Any (large) theatre with 1 more person than an empty theatre is empty.

Conclusion. A theatre with 500 people in it is empty.

Burnett attributes these types of vagueness effects to an absolute adjective's context-sensitive *tolerant* extension.

While relative adjectives obey a version of van Benthem's axioms (see 28 above), absolute adjectives have invariant denotations across comparison classes (48).[28]

(48) **Absolute Adjective Axiom**: For all absolute predicates Q, all interpretations $[\![\cdot]\!]^M$, all comparison classes $X \subseteq D$, and all individuals $a \in X$,

a. if $[\![Q(a)]\!]^{M,X} = 1$, then $[\![Q(a)]\!]^{M,D} = 1$, and

b. if $[\![Q(a)]\!]^{M,D} = 1$, then $[\![Q(a)]\!]^{M,X} = 1$

The combination of (46b) and (48) has the effect of rendering the strict denotation of absolute adjectives equivalent to their classical denotation, but does not affect their tolerant denotation. For example, take a T-model with $D = \{a, b, c\}$, a in the classical denotation of an absolute adjective Q, and an indifference relation defined as in (49).

27. As Burnett explains, these symbols are meant to be read "b is relevantly indifferent to a" and "a is relevantly indifferent to b", respectively.

28. I have slightly altered Burnett's presentation.

(49) $\{a \sim a, b \sim b, c \sim c, b \sim a\}$

Since a is the only member of Q's classical denotation, and since for all d such that $a \sim d$, $[\![P(d)]\!]^c = 1$ (the only such d is a itself, given the directionality of Burnett's indifference relation $\sim$), $P(a)$ (and only $P(a)$) is s-true in any comparison class composed of (some combination of) $\{a, b, c\}$. In contrast, $[\![P(b)]\!]^t = 1$ in the comparison class $\{a, b, c\}$, but $[\![P(b)]\!]^t = 0$ in the comparison class composed of $\{b, c\}$.[29] Thus, the tolerant denotation of an absolute predicate varies, but the strict denotation remains constant. The variable tolerant denotation accounts for vagueness effects while the stable strict denotation accounts for the absolute adjective's seeming lack of vagueness effects.

Many-valued logic's choice rule

Cobreros et al. (2012) have built intransitivity of indifference into their T-models. Burnett elaborates their system to include a semantics for comparatives, including tolerant comparison and strict comparison (50).

(50) a. $[\![\text{a} > \text{b}]\!]^c$ =1 iff there is some $X \subseteq D$ such that $[\![\text{P(a)}]\!]^{c,X}$ =1 and $[\![\text{P(b)}]\!]^{c,X}$ =0

b. $[\![\text{a} > \text{b}]\!]^t$ =1 iff there is some $X \subseteq D$ such that $[\![\text{P(a)}]\!]^{t,X}$ =1 and $[\![\text{P(b)}]\!]^{t,X}$ =0

c. $[\![\text{a} > \text{b}]\!]^s$ =1 iff there is some $X \subseteq D$ such that $[\![\text{P(a)}]\!]^{s,X}$ =1 and $[\![\text{P(b)}]\!]^{s,X}$ =0

The classical comparative in (50a) is equivalent to the method Klein (1980) used to define comparatives, and therefore this comparative gives rise to a transitive comparative relation. As Burnett argues, the tolerant comparative relation also gives rise to a transitive comparative relation for absolute adjectives.

Yet while the tolerant-classical-strict setup has the potential to generate the right type of choice rule, the setup is too rigid to model higher-order vagueness, as Cobreros et al. acknowledge. As it currently stands, the indifference relation itself has no borderline cases.

29. As this example demonstrates, tolerant truth fails to respect consistency under set contraction, Sen's principle α.

Furthermore, while tolerant truth provides for great flexibility—including validating the sorites premise—this flexibility comes at the cost validating certain contradictions (see Alxatib et al. (2013) for pragmatic problems with Cobreros et al.'s proposal). But while (51a) may be acceptable to some speakers, (51b) seems much more degraded. Cobreros et al. cannot explain why this should be.

(51) a. Clarence is tall and Clarence is not tall.

b. ?? Clarence is tall in one sense and Clarence is not tall in the same sense.

The natural counterpart to (51b) is *Clarence is tall in one sense and Clarence is not tall in another sense*, and I will argue that this is the analysis (51a) should take. Therefore, the validation of contradictions like (51b) remains a serious drawback.[30]

Finally, Burnett's implementation of the tolerant-classical-strict framework creates its own problems. First, Burnett asserts that the axioms constraining the indifference relation are *pragmatic* in nature. It is hard to make sense of this claim in her system, though, because tolerant truth and strict truth are defined in terms of the indifference relation, and the pairs related by the indifference relation are stipulated as part of the T-model. Second, Burnett's tolerant comparative relation (50b) gives rise to a comparative relation that allows a to be P-er than b in one context but not in another. Let $D = \{a, b, c\}$, with $[\![P(a)]\!]^c$ and $b \sim a$. Then in $X = \{a, b, c\}$, $[\![P(b)]\!]^t = 1$ and $[\![P(c)]\!]^t = 0$. Therefore, $[\![b > c]\!]^t = 1$. However, in the context $X = \{b, c\}$, $[\![b > c]\!]^t = 0$. The truth of *a is taller than b* does not vary depending on entities other than a and b.[31]

2.3.4 Contextualism

Contextualist approaches to vagueness vary widely. As Smith (2008, 113) describes contextualist accounts, "The basic idea behind contextualism is that vagueness is a diachronic

30. I take Cobreros et al.'s intuitions that certain contradictions are valid seriously, but I also take seriously Williamson's rejoinder: it is difficult to argue with those who contradict themselves.

31. But see the discussion of Independence of Irrelevant Alternatives in 3.

phenomenon, which only emerges when we consider the semantic state of a language over time (or more generally, over multiple instances of interpretation)." More particularly, as Fara (2002); Raffman (1994); Shapiro (2006) all agree, contextualist accounts place particular emphasis on the act of interpretation, arguing that the act itself changes the semantic "facts on the ground."

For some contextualists, like Kamp (1981), the idea is that the context or valuation function changes as we proceed along the steps in the sorites argument (Raffman, 1994, 1996; Fara, 2002; Shapiro, 2006). Kamp's particular argument involved evaluating the consequent of a conditional in a context modified by taking the antecedent as true. That is, for the conditional $(\phi(x_n) \rightarrow \phi(x_{n+1}))$, $\phi(x_n)$ is evaluated in a different context than $\phi(x_{n+1})$. The new valuation function, importantly, requires that entities indistinguishable from x_n have property ϕ. (A conditional is true in a context c just in case its antecedent is false or its consequent *would* be true in the context that results from taking $\phi(n)$ to be true, as Fara (2002) paraphrases.) Therefore, the sorites premise is false but appears to be true because the argument relies on an equivocation with respect to the context of interpretation.

Adopting a slightly different approach, Fara (2002, 59) relies on the Similarity Constraint in (52) to derive vagueness effects.

(52) **Similarity Constraint.** Whatever standard is in use for a vague expression, anything that is saliently similar, in the relevant respect, to something that meets the standard itself meets the standards; anything saliently similar to something that fails to meet the standard itself fails to meet the standard.

The way the Similarity Constraint works in practice, as Fara explains, is the following: "for any particular x and y that differ in height by just 1mm, the very act of our evaluation raises the similarity of the pair to salience, which has the effect of rendering *true* the very instance we are considering." Like Kamp, Fara affirms each instance of the sorites premise (*if* n *is* P*, then* $n+1$ *is* P) without affirming the universally-quantified version ($\forall n$ *if* n *is* P*, then* $n+1$ *is* P). The sorites premise is therefore false.

If we consider each instance of the sorites premise simultaneously, however, the sorites paradox resurfaces. Fara's response is that such simultaneous consideration is impossible. In other words, each comparison cannot be simultaneously *salient* to the speaker: "there are too many pairs for us to actively entertain each similarity." Moreover, Fara's notion of salience incorporates our expectations, interests, etc. Depending on these interests, Fara explains, the utterance *Elena is tall* means that there is a difference between Elena's height and the height norm (for some comparison class), and that this difference is *significant* (53).

(53) $[\![\text{Elena is tall}]\!] = \text{tall}(e) \gg \text{NORM}(\text{tall})$

Extending a similar contextualist approach to model the difference between absolute and relative adjectives, van Rooij (2011) builds on the underlying framework of Klein (1980). This framework holds that the meaning of the positive form is relative to a context-dependent comparison class. Van Rooij defines a *context structure* $\langle I, C, V \rangle$ with I a nonempty set of individuals, C the set of finite subsets of I, and a valuation V that assigns to each $c \in C$ those individuals in c that count as being P, for some predicate P. Then, $P(c) = \{x \in c \mid x \in V(P, c)\}$. The idea, as van Rooij, 140 explains (and as I have already noted), is that "P is thought of as a choice function, selecting the *best* elements of c."[32]

32. Van Rooij's initial account of absolute adjectives seems to falter. He explains that we can constrain the behavior of adjectives in diverse comparison classes by subjecting them to certain consistency axioms. His chosen axioms derive from Arrow (1959), who showed how to generate an ordering relation from the behavior of a choice function over different sets. Analogizing this behavior to adjectival denotations in different comparison classes, van Rooij attempts to adapt Arrow's (1).

(1) If $c \subseteq c'$ and $c \cap P(c') \neq \emptyset$, then $P(c) = P(c') \cap c$.

Arrow claims that (1) has the intuitive interpretation that, if some elements are chosen from c', and if the range of alternatives is reduced to c with some previously chosen elements, then no previously unchosen element becomes chosen and no previously chosen element becomes unchosen.

Van Rooij decomposes (1) into two separate axioms.

(2) a. If $c \subseteq c'$, then $P(c') \cap c \subseteq P(c)$.

b. If $c \subseteq c'$ and $c \cap P(c') \neq \emptyset$, then $P(c) \subseteq P(c')$.

Van Rooij then imposes conditions on this choice function that are meant to generate vagueness effects. These conditions amount to a semi-order. As explained in connection to epistemicism, the idea behind a semi-order is that it introduces some threshold ϵ within which entities may be related by the indifference relation $\sim$. A semi-order validates the tolerance principle because, for any two adjacent items in a sorites series, they will be related by the indifference relation. But items far enough apart in the sorites sequence are *not* indifferent with respect to P.

Combining the notion of a semi-order with Fara-style contextualism, van Rooij relies on the notion of a *pragmatic comparison class* to solve the problem presented by considering each instance of the sorites premise simultaneously . He claims that a vague predicate P is used appropriately only if it is used in a context in which it clearly demarcates the P entities from the non-P entities. As with Fara's explanation, the tolerance principle is true for two similar entities when actively considered, but no contradiction follows because it is "pragmatically inappropriate" to confront a "set of objects that form a Sorites series" (van Rooij, 2011, 155).

He then attempts to explain the data in (3).

(3) The desk is not flat, but it is flatter than the pavement.

Since *flat* is an absolute adjective, van Rooij requires that its interpretation is relative to the entire domain ("in the positive use of 'flat', the adjective should always be interpreted with respect to the *maximal* comparison class: the whole domain" (van Rooij, 2011, 139)). The comparative *the desk is flatter than the pavement* in (3) is subject to two constraints, (i) $\{d, p\} \subseteq c$ and (ii) $d \in \text{flat}(\{d, p\})$ but $p \notin \text{flat}(\{d, p\})$. Van Rooij then concludes that (3) does not denote a contradiction because "the desk can be flat compared to the pavement, without it being the case that the desk is flat when compared to other members of $c \supseteq \{d, p\}$" (van Rooij, 2011, 141).

But that is not quite right. Van Rooij claims that, by (2a), $\{d, p\} \cap \text{flat}(c) \subseteq \{d\}$. If that's the case, though, then $\{d\} \subseteq \text{flat}(c)$, which means that the desk (d) is flat compared to the superset comparison class c. It is therefore *not* the case that "the desk can be flat compared to the pavement, without it being the case that the desk is flat when compared to other members of c."

Contextualism's choice rule

Because contextualism is built on top of other theories of predicate meaning, contextualist accounts generally inherit the choice rule of their substrate. For example, Fara presents a contextualist solution to the sorites paradox, but she relies on an epistemicist semantics: predicates have boundaries between their positive and negative extensions. The choice rule to be inferred from her account is the same as that inferred from Kennedy's (2007). The contextualist account of van Rooij, in contrast, relies on a supervaluationist substrate and therefore inherits the choice rule generated by conditions imposed on the relations between comparison classes.

In addition, however, these accounts may suffer from an additional drawback identified by Stanley (2003). If, as Stanley argues, the content of a vague predicate must be invariant under VP ellipsis, then sorites paradoxes like (54) are immune to the contextualist's analysis.

(54) If that_1 is a heap, then that_2 is too; if that_2 is a heap, then that_3 is too, ..., if that_{n-1} is a heap, then that_n is too.

As Stanley (2003, 272) explains, "Since no context-shifts are possible in these cases, the contextualist's semantic account fails to explain why we find each step in [(54)] so compelling."[33] Finally, contextualist accounts are difficult to square with non-adjectival sources of vagueness, as Fara acknowledges.

2.4 Conclusion

On a choice-functional approach to predicate meanings, the denotation of a vague predicate like *tall* looks like (55), where the choice rule for tallness is to the right of the equivalence.

(55) $[\![\text{tall}]\!]^c = \{x \mid \forall y \in c,\ xRy\}$

The linguistic phenomena associated with vagueness—the sorites paradox, borderline cases,

33. But see Fara (2008).

absolute and relative adjectives—tell us that the relation R associated with healthiness must have the property of being intransitive in its indifference relation. The linguistic phenomena associated with comparatives tell us that this relation must have the property of being transitive in its comparative relation. In general, the theories of vagueness discussed in this chapter have tried to respect these requirements. In the process, though, each of these theories has generated its own problems, either through the abandonment of classical logic (supervaluationism, many-valued logic, contextualism) or through a deeply counterintuitive relationship between use and meaning (epistemicism).

These theories share choice rules but differ greatly on how such rules affect the truth of propositions like *Sonia is tall*. This suggests that a theory of vagueness should start at the choice rule—the domain of the (55) rather than the range. In the next chapter, I begin this project, showing how the shared choice rule is natural in a multidimensional context, where choice is attentive to multiple dimensions.

CHAPTER 3

SEMANTIC INDECISION

"Vagueness is semantic indecision," David Lewis (1986, 213) said. If vagueness is a type of decision problem, then the sorites paradox is a paradox of decision making. Vague predicate result from the the aggregation of many judgments into one, and this aggregation procedure is subject to paradoxes of collective decision making. In 1785, the Marquis de Condorcet noticed a curious feature of elections to the French Academy of Sciences. If three equal blocks of voters had certain preferences among the available candidates (as in (56)), then majority vote would fail to deliver an outcome.

(56)	Block 1:	Candidate A > Candidate B > Candidate C
	Block 2:	Candidate C > Candidate A > Candidate B
	Block 3:	Candidate B > Candidate C > Candidate A
	Election:	Candidate A > Candidate B > Candidate C > Candidate A

Condorcet's paradox and the sorites paradox are two sides of the same coin. They both arise from a common source, the aggregation of many judgments into one. And in both cases, plausible constraints on the aggregation process lead to paradox.

In order to make the argument that vagueness effects arise from paradoxes of collective choice, I draw on the tools of social choice theory (the branch of economics devoted to collective decision making) to show how using a vague predicate is like a decision problem. I begin with evidence from "multidimensional" adjectives, like *healthy.* These adjectives aggregate measures along multiple dimensions into one scale. I reinterpret Arrow's Theorem—a generalization of Condorcet's paradox—in the choice space of multidimensional adjectives in two ways. First, reasonable semantic constraints on what multidimensional adjectives should be like lead to paradox. Second, this paradox arises from what Saari (2008) calls the "curse of dimensionality," the multidimensionality of the choice space itself. Vagueness is semantic indecision because certain choices in multidimensional space are impossible.

3.1 Limitations on collective choice

Collective choice functions are subject to important limitations. The Marquis de Condorcet observed that cycles (or intransitivities, e.g. A is preferred to B is preferred to C is preferred to A) may arise in a collective body like a legislature whenever such a body tries to choose among three or more proposals. In (57), for example, pairwise majority vote delivers the election of Candidate A over Candidate B (Blocks 1 and 2) and Candidate B over Candidate C (Blocks 1 and 3). By similar reasoning, Candidate C beats Candidate A (Blocks 2 and 3). As a whole, the legislature has cycling preferences: it prefers A to B to C to A.

(57) Block 1: Candidate A > Candidate B > Candidate C

Block 2: Candidate C > Candidate A > Candidate B

Block 3: Candidate B > Candidate C > Candidate A

Election: Candidate A > Candidate B > Candidate C > Candidate A

Kenneth Arrow showed that there is no collective decision procedure that respects certain reasonable assumptions and avoids intransitivity (Arrow, 1950). Let $X = \{x, y, z, \ldots\}$ be candidates in a voting procedure and let $N = \{1, \ldots, n\}$ be the set of voters. By hypothesis, these voters have transitive preferences. Arrow's assumptions are set out informally in (58).

(58) a. Unanimity: If every voter prefers alternative x to y, the collective choice prefers x to y.

b. Independence of Irrelevant Alternatives: The collective choice between x and y depends only how the voters rank x and y.

c. Unrestricted Domain: The voters can rank the relevant candidates in any way.

d. Non-Dictatorship: No one voter determines the result of the collective choice.

(59) Arrow's Theorem: There is no complete and transitive collective decision procedure f that respects all four constraints.

If Arrow's constraints are obeyed (and they have some normative appeal in the voting context), the collective decision procedure will be incomplete or intransitive (or both) (164).

More specifically, Arrow's theorem holds that for a finite number of voters and at least three distinct alternatives, there is no social welfare function f that is both transitive and complete. I provide a sketch of a proof of Arrow's Theorem following the strategy presented in Gaertner (2009).

Within an interpretive context, fix a finite set $X = a, b, c, \ldots, n$ of alternatives and a finite set $D = 1, 2, 3, \ldots, i$ of voters. Let $\succsim_i$ be the ordering of voter i's preferences with respect to the alternatives. Thus $a \succsim_i b$ means "voter i prefers alternative a at least as much as alternative b." Any voter's preference ordering is transitive and complete.

(60) Transitive: If $a \succsim_i b$ and $b \succsim_i c$ then $a \succsim_i c$, for all a, b, c in X and for all i.

(61) Complete: For any $\succsim_i$ and any a, b, either $a \succsim_i b$ or $b \succsim_i a$.

Let the relation $\succsim_{\mathrm{S}}$ represent the "social ordering" derived from aggregating the set of voter preferences $\succsim_1, \succsim_2, \ldots, \succsim_i$. A "social welfare function" is a mapping f from the set of voters' preferences to the social ordering (62). Any particular set of voters' preferences $\succsim_1, \succsim_2, \ldots, \succsim_i$ is called a "profile."

(62) $\succsim_{\mathrm{S}} = f(\succsim_1, \succsim_2, \ldots, \succsim_i)$

The Arrowian axioms are presented in formal terms below.

(63) a. Unanimity: For all i and every pair of alternatives a, b in X, if $a \succsim_i b$ for all voters i, then $a \succsim_{\mathrm{S}} b$.

b. Independence of Irrelevant Alternatives: For every pair of entities a and b and every pair of "profiles" $(\succsim_1, \succsim_2, \ldots, \succsim_i), (\succsim_{1'}, \succsim_{2'}, \ldots, \succsim_{i'})$, if for all voters i the relations $\succsim_i$ and $\succsim_{i'}$ coincide on a, b, then $\succsim_{\mathrm{S}}$ and $\succsim_{\mathrm{S}'}$ coincide on a, b. By "coincide" I mean that both relations rank a over b, both rank b over a, or both rank a and b equally.

c. Unrestricted Domain: The domain of the social welfare function f potentially includes every list $\langle \succsim_1, \succsim_2, \ldots, \succsim_i \rangle$ of i weak orderings of the alternatives in X.

d. Non-Dictatorship: There is no voter i such that $f(\succsim_1, \succsim_2, \ldots, \succsim_i) = \succsim_i$ for all profiles $(\succsim_1, \succsim_2, \ldots, \succsim_i)$.

First, we define what it is for a voter or a set of voters to be "almost decisive" and "decisive" among two alternatives.

(64) Almost Decisive. A set of voters V is almost decisive for some a against some b if, whenever $a \succ_{\mathrm{i}} b$ for every i in V and $b \succ_{\mathrm{j}} a$ for every j outside of V, $a \succ_{\mathrm{S}} b$.

(65) Decisive. A set of voters V is decisive for some a against some b if, whenever $a \succ_{\mathrm{i}} b$ for every i in V, $a \succ_{\mathrm{S}} b$.

Let $D(a, b)$ stand for the proposition that a voter i is "almost" decisive with respect to a and b, and $\overline{D}(a, b)$ for the proposition that i is decisive with respect to a and b.

This allows us to prove the following lemma:

(66) Lemma. If there is some voter i that is almost decisive for some pair of alternatives (a, b), an Arrowian social welfare function f satisfying the first three axioms (Unanimity, Independence, and Unrestricted Domain) implies that i is a dictator.

To prove this, assume that i is almost decisive with respect to a and b, that is $D(a, b)$. For a third alternative c and for all $n \neq i$, let

(67) $a \succ_{\mathrm{i}} b, b \succ_{\mathrm{i}} c$

$b \succ_{\mathrm{n}} a, b \succ_{\mathrm{n}} c$

By $D(a, b)$, we obtain $a \succ_{\mathrm{S}} b$. By the Unanimity condition, we obtain $b \succ_{\mathrm{S}} c$. And by the (presumed) transitivity of f, we obtain $a \succ_{\mathrm{S}} c$. Note that the ordering with respect to a and c according to voter i was left unspecified. But this means that $a \succ_{\mathrm{S}} c$ is a consequence of $a \succ_{\mathrm{i}} c$ alone. And this, in turn, means that $D(a, b) \rightarrow \overline{D}(a, c)$. The property of decisiveness is "catching": whereas we only assumed i was almost decisive with respect to a and b, we have discovered that this assumption means that i is decisive with respect to a and c.

So we try a different permutation of the logical possibilities for the profiles of the voters. Suppose, as before, that $D(a, b)$, but now the individual preferences are as follows:

(68) $c \succ_{\mathrm{i}} a, a \succ_{\mathrm{i}} b$

$c \succ_{\mathrm{n}} a, b \succ_{\mathrm{n}} a$

As before, the aggregated ranking is $a \succ_S b$. And due to Unanimity, we obtain $c \succ_S a$. Therefore, the aggregated preference is $c \succ_S b$ by transitivity. But the only dimension that ranks c above b is i. This shows that $D(a, b) \rightarrow \overline{D}(c, b)$.

In fact, for any permutation of voter preferences among a, b, and c (with a similar argument holding for any finite number of alternatives), it is possible to show that, if i is almost decisive as to one pair, i is decisive for every pair (Gaertner, 2009). That is, i is a dictator, contrary to the Non-Dictatorship axiom.

Since the first three axioms imply that i is a dictator if i is almost decisive with respect to any two alternatives, the axiom of Non-Dictatorship requires that we cannot allow i to be almost decisive for any two alternatives. But obeying this requirement leads to a contradiction.

The Unanimity condition entails that there is at least one decisive set for any ordered pair (a, b): the set of all voters. Since "decisive" entails "almost decisive," there is also at least one almost decisive set of voters for some pair of alternatives. Choose the smallest such set. This set must contain at least two voter since, as shown by the Lemma above, if only one voter is almost decisive, that voter is a dictator. Call the smallest almost decisive set of voters V and let V be almost decisive for (a, b).

Now we divide V into two parts: V_1 contains one voter, and V_2 contains the rest of V. Suppose that the voters have the following profiles:

(69) For i in V_1: $a \succ_i b$ and $b \succ_i c$

For all n in V_2: $a \succ_n b$ and $c \succ_n a$

Because V is almost decisive for a and b, we obtain $a \succ_S b$. Now we are faced with a dilemma: if $c \succ_S b$ holds, then V_2 is almost decisive for (c, b) (this is true even assuming all the voters outside of V rank b over c). But since we assumed that V is the smallest decisive set, and since V_2 is a strict subset of V, V_2 cannot be almost decisive. So we must have *not* $c \succ_S b$. This means that we have $b \succsim_S c$. By transitivity of the aggregated ordering, $a \succ_S c$. But only i has this ordering. This means that V_1 is almost decisive, again contradicting our assumption that V is the smallest decisive set.

This concludes the proof: the aggregation of voter preferences cannot obey all four axioms and still be transitive and complete.

In addition to intransitivity, a collective decision procedure f may display discontinuities. That is, small changes in the voters' preferences do not necessarily result in a small change in the outcome (Gaertner, 2009). For example, Gaertner discusses the following example: consider voters 1-5 with preferences among the alternatives v, x, y, z in descending order, as in (70).

(70)

voter 1	voter 2	voter 3	voter 4	voter 5
x	y	z	x	z
y	v	v	v	x
z	x	y	z	v
v	z	x	y	y

Alternative x is the "Condorcet winner" in this case because it beats all the other alternatives by majority vote in a pairwise contest. However, if the order of x and z are reversed in voter 2's preference ranking—voter 2's least prefered alternatives—then z is the Condorcet winner. A small change in the voters' preferences results in a large change in the outcome to the election.

This behavior was generalized by Chichilnisky (1982). Suppose two people, n and e want to go camping along the shore of a perfectly circular lake. (This example is presented by Saari (1997).) They may prefer the same geographic location along the shore, or they may not. If they agree on the location, that is where they will camp. This is an analogue to Arrow's Unanimity.

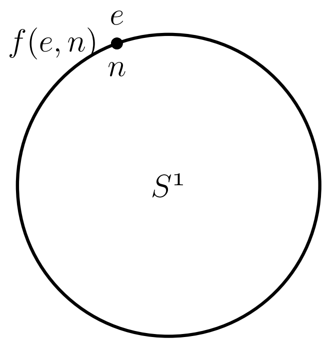

(71)

And the decision will be anonymous, meaning that it will not depend on who chose what. This is an analogue to Arrow's Non-dictatorship.

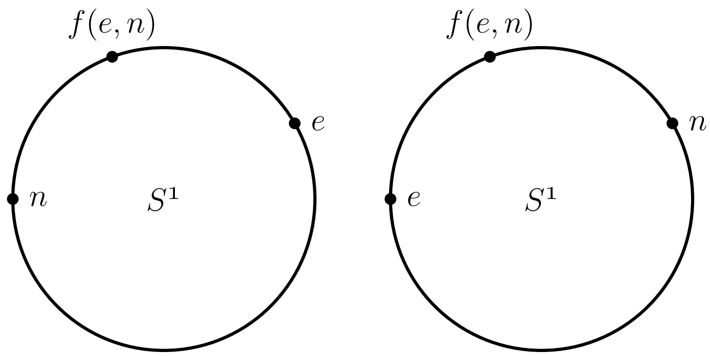

(72)

Finally, the decision rule should be "relatively insensitive to small changes in individual preference" (Chichilnisky, 1982, 337). This last requirement is continuity.[1]

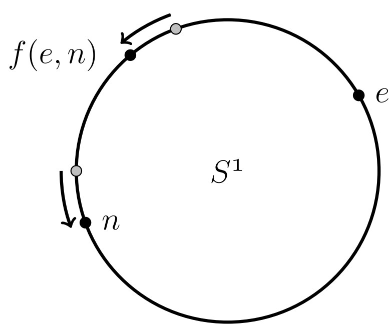

(73)

Chichilnisky showed that there is no unanimous, anonymous, and continuous collective choice function f (74). To get a feel for why, hold e fixed. As n moves continuously in a counter-clockwise direction, so does $f(e, n)$, until n reaches the antipode of e. Then, $f(e, n)$ jumps to the other side of the lake.

1. Baryshnikov (1997) emphasizes the relationship between Independence of Irrelevant Alternatives and continuity.

Figure 3.1: Aggregation of location preferences along a circular lake (Grinsell, 2014).

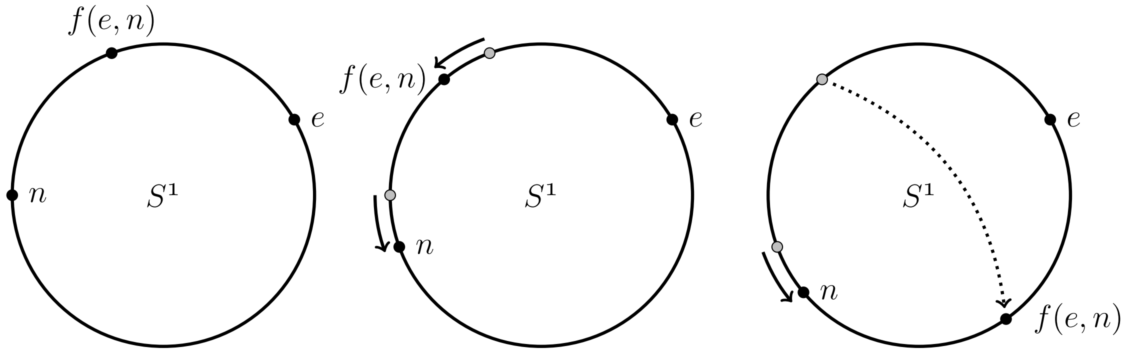

(74) Chichilnisky's Theorem: There is no continuous aggregation rule $f : S^1 \times S^1 \to S^1$ that satisfies unanimity and anonymity.

As discussed below, these two limitations on collective choice, intransitivity and discontinuity, are both reflected in vagueness phenomena.

3.2 Choice theory for adjectival semantics

The path connecting limitations on collective choice as represented by Arrow's theorem and vagueness effects is visible on a choice-theoretic interpretation of adjectival semantics. As discussed in 2, a choice-theoretic approach to adjectival semantics is familiar, road-tested, and easy to reconcile with more traditional theories of adjectival semantics.

Choice functions are defined as follows. Let X be a set of available options. Define a *choice structure* as in (75, see Mas-Colell et al. (1995)):[2]

(75) $(\mathcal{B}, C(\cdot))$, where

2. This is analogous to Reinhart's (1997: 372) definition:

(1) A function f is a choice function ($CH(f)$) if it applies to any non-empty set and yields a member of that set.

However, it is not equivalent. In particular, Reinhart is not concerned with the type of choice rules discussed below.

a. $\mathcal{B}$ is a family of nonempty subsets of X; that is, every element of $\mathcal{B}$ is a set $B \subset X$, and

b. $C(\cdot)$ is a choice rule that assigns a nonempty set of chosen elements for every set B; that is, $C(B) \subset B$.

An example of a choice structure is (76), repeated from (17).

(76) $(\mathcal{B}, C_1(\cdot))$, where $X = \{x, y, z\}$ and $\mathcal{B} = \{\{x, y\}, \{x, y, z\}\}$. The choice rule $C_1(\cdot)$ may be defined as $C_1(\{x, y\}) = \{x\}$ and $C_1(\{x, y, z\}) = \{x, y\}$.

First, choice functions are already a familiar tool in the semantics toolbox. Perhaps their most conspicuous use appears in the analysis of indefinites, free choice items, and quantifiers (see Reinhart (1997); Winter (1997); Giannakidou (2004) and many others). But they also appear in modal semantics (see, e.g., Kratzer (1981)), the semantics of focus (Reich, 2004), and other areas.

Second, a choice-theoretic approach is straightforwardly extendable to an analysis of adjectival semantics, and van Rooij's (2011) account proceeds along these lines. Building on Klein's (1980) framework, van Rooij treats gradable adjectives as choice functions, selecting entities from the comparison class (the set of entities somehow relevant to the interpretation of the gradable adjective).[3] The gradable adjective P can be "thought of as a choice function, selecting the *best* elements of [the comparison class] c" (van Rooij, 2011, 140).

If the comparison class forms the set of available options, the next question is how to model the choice rule among these options. Van Rooij imposes conditions that mimic the requirements of the Weak Axiom of Revealed Preference (18). For example, he wants to rule out the case in which a counts as P in the comparison class $\{a, b, c\}$ but a does not count as P in the comparison class $\{a, b\}$.[4] By imposing conditions like (18), van Rooij

3. For example, a sentence like *Ruth is tall* appears to make little sense unless we are comparing Ruth's height to the height of relevant individuals (Klein, 1980). The set of relevant individuals is the comparison class.

4. As noted in 2, this condition and others like it are analogous to Klein's (1980) "Consistency Postulate"

generates a weak order among the P entities, thereby deriving the "penumbral connections" that account for the behavior of gradable adjectives across comparison classes (Fine, 1975).[5]

Third, familiar degree-theoretic accounts can be recast in choice theoretic terms. On Kennedy's (1999) degree-theoretic account, for example, gradable adjectives like *tall* are measure functions that map an entity to a position on a one-dimensional scale of height, as in Fig. 3.2 (Kennedy, 1999). These measure functions combine with a phonologically null morpheme to derive the denotation in (77).

(77) $[\![\text{tall}]\!] = \lambda x.\text{tall}(x) \geq \mathbf{s}(\text{tall})$

Loosely paraphrased, x is tall if x's height is significantly greater than some "standard of comparison," which represents the cutoff point between positive and negative extensions.[6] The choice rule associated with this type of "significant difference" approach accounts for behavior across comparison classes by generating a semi-order among the P entities (Luce, 1956).[7]

---

(1).

(1) For entities a, b and predicate P, if a counts as P in comparison class c, and if $b \geq a$ with respect to the ordering of the domain of P, then b counts as P in c. Similarly, for a comparison class c' in which a counts as $\neg P$ and if $a \geq b$, then b counts as $\neg P$ in c'.

Van Rooij imposes the following axioms on the behavior of the choice function.

(2) a. If $c \subseteq c'$, then $P(c') \cap c \subseteq P(c)$.

b. If $c \subseteq c'$ and $c \cap P(c') \neq \emptyset$, then $P(c) \subseteq P(c')$.

These axioms (and one other) also generate a weak order.

5. Van Rooij ultimately settles on a semi-order to represent vagueness effects, see section 2.3.4.

6. More precisely, the standard-setting function $\mathbf{s}$ "is a context-sensitive function that chooses a standard of comparison in such a way as to ensure that the objects that the positive form is true of 'stand out' in the context of utterance, relative to the kind of measurement that the adjective encodes" (Kennedy, 2007, 17).

7. The ability to generate a choice out of some set and the ability to impose some sort of ordering on that set are closely connected (Mas-Colell et al. 1995: §1D). In what follows, I sometimes talk of choices and sometimes of orderings or rankings. For my purposes, there is no harm in this alternation.

Figure 3.2: Measure functions.

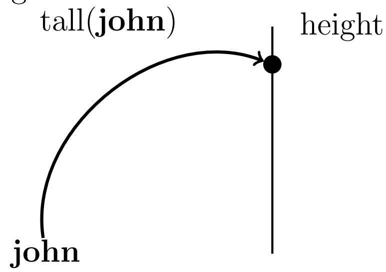

A choice-theoretic account is also straightforwardly extendable to comparatives. As discussed in chapter 57, transitivity is built into to the semantics of comparatives, and theories of vagueness must negotiate this requirement. In traditional degree semantics (see, e.g., Cresswell (1976b); von Stechow (1984); Heim (1985, 2001); Kennedy (1999)), gradable predicates like *tall* or *healthy* map individuals onto scales (as in Fig. 3.2). Expressing this same idea in terms of relations, gradable predicates denote relations between individuals and degrees, as in (78).

(78) $[\![\text{tall}]\!] = \lambda d_d \lambda x_e.\textsc{height}(x) \geq d$

(or $\lambda d \lambda x.\text{tall}(x, d)$)

Comparatives are relations between those degrees.

In a comparative like (79), the comparative morpheme *-er* (80) compares the maximal degree to Sam is tall to the maximal degree to which John is tall (Fig. 3.3).

(79) Sam is taller than John.

(80) $[\![\text{-er}]\!] = \lambda P_{d,t} \lambda Q_{d,t}.\max(Q) > \max(P)$

This derives the correct meaning for a comparative like (79).

(81) a. $\lambda d.[\textsc{height}(\text{Sam}) \geq d]$ = degrees of Sam's tallness

b. $[\![\text{-er}]\!]$ ($[\![\text{Sam is tall}]\!]$) ($[\![\text{John is tall}]\!]$)

$[\lambda P_{d,t} \lambda Q_{d,t}.\max(Q) > \max(P)](\lambda d'.[\textsc{height}(\text{John}) \geq d'])(\lambda d.[\textsc{height}(\text{Sam}) \geq d])$

Figure 3.3: Comparison

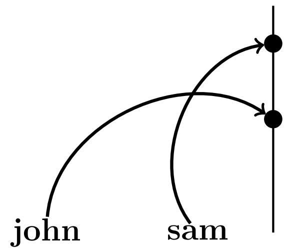

c. $\max(\lambda d'.[\textsc{height}(\text{Sam}) \geq d']) > \max(\lambda d.[\textsc{height}(\text{John}) \geq d])$

'the maximum degree to which Sam is tall is greater than the maximum degree to which John is tall'

The semantics of the comparative morpheme (80) takes advantage of the ordering relation $>$, typically defined over a linear scale. On any such scale, the relation $>$ will be transitive by definition. As suggested in chapter (2), on a choice-theoretic interpretation of adjectival semantics, this relation is the counterpart of the strict preference relation.

3.3 Multidimensional adjectives

Multidimensional adjectives complicate the relationship between comparison and transitivity. However, they fit nicely into a choice-theoretic view of adjectival semantics once they are understood as *collective* choice functions—the type of choice functions subject to Arrow's result.

Multidimensional adjectives are associated with multiple dimensions (or criteria) of evaluation (Kamp, 1975; Klein, 1980; Sassoon, 2013a,b).[8] For example, *similar, identical, typical, normal, good, happy,* and *healthy* are all multidimensional adjectives. Whether Clarence is healthy intuitively depends on a number of factors simultaneously, like blood pressure, heart function, cholesterol levels, and more. Context normally supplies the relevant dimensions.

Sassoon (2013a) provides a number of ways to identify multidimensional adjectives. For

8. The sense in which "dimension" is used here is distinct from another use in semantic theory meant to distinguish between at-issue and implicated meaning (Potts, 2005).

example, it is possible to specify a dimension overtly with prepositional phrases like *with respect to* or *in*, as in (82). Such specifications are not available for "unidimensional" adjectives like *long*. Despite having both temporal and spatial interpretations, a single utterance of *long* is not interpreted relative to both dimensions at the same time(82b).

(82) a. John is healthy with respect to blood pressure.

b. * The wedding is long {with respect to, in} temporal duration (but not with respect to space).

It is also possible to quantify over multiple dimensions, as in (83), and to target dimensions with *wh*-words (84).

(83) a. Elena is healthy in {every respect, some respects, most respects}.

b. * The table is long in {all, most, three, some} respects.

(84) In what respect is Elena healthy?

And it is possible to except dimensions from consideration, as in (85).

(85) a. Ruth is healthy except with respect to her cholesterol.

b. * The table is long except with respect to temporal duration.

These tests suggest that an adjective like *healthy* is semantically sensitive to multiple dimensions.

Moreover, some multidimensional adjectives pattern with absolute adjectives, as in (86a), but still display vagueness effects. An adjective like *healthy* is compatible with the modifier *perfectly*, which picks out the maximal element in a set of degrees. It is not (or at least not as) compatible with the modifier *slightly*, which requires a minimal element in a set of degrees. This is the same pattern that upper closed adjectives like *straight* display. The multidimensional adjective *abnormal* follows the same patterns as lower closed adjectives like *bent* (see 86b). Still other multidimensional adjectives like *familiar* follow the totally closed patterns, as in (86c).

(86) a. Stephen is {perfectly, almost, ??slightly} healthy.

b. Tony is {??perfectly, slightly} abnormal.

c. Ruth is {perfectly, slightly} familiar with these routines.

These multidimensional adjectives also behave like absolute adjectives in the definite description test. Imagine two sickly looking plants, one much more sickly looking than the other. My intuition is to reject (87).

(87) ?? Pass me the healthy one. Syrett et al. (2010)

But if these tests suggest an endpoint-oriented interpretation, as with absolute adjectives, they do not signal a lack of vagueness effects. In particular, multidimensional adjectives like *healthy* behave in a way that is inconsistent with Interpretive Economy. For instance, a totally closed multidimensional adjective like *familiar* may have a non-endpoint-oriented interpretation, as the parenthetical followup to (88) demonstrates (McNally, 2011).

(88) For a new Justice, Elena is familiar with the Court's routines. (In fact, she's completely familiar with them.)

More importantly, endpoint-oriented multidimensional adjectives display vagueness effects like borderline cases (as in 89) and Sorites premises (as in 90).

(89) A person with low blood pressure, low cholesterol, and a low heart rate is healthy. A person with high blood pressure, high cholesterol, and a high heart rate is not healthy. Is a person with low blood pressure, low cholesterol, and a high heart rate healthy? A person with mid-values for all the measures?[9]

(90) A person whose systolic blood pressure is one unit higher than a healthy person's is healthy.

9. Sassoon claims that *healthy* means something like 'healthy in all respects,' so an unhealthy measure on one dimension disqualifies the entity from being *healthy.* My own intuitions are not robust here, though even if Sassoon is correct, there are still borderline cases of *healthy* (e.g. the person with mid-values for all the measures).

And multidimensional adjectives like *healthy* may have seemingly intransitive comparisons (91).

(91) John is healthier than Ruth.
Ruth is healthier than Sam.
Sam is healthier than John.

The comparatives in (91) are "seemingly" intransitive because (91) involves comparison along implicitly different dimensions. This is brought out by (92), which makes explicit the relevant dimension. Once the dimension is fixed across the three comparatives, the comparative relation is no longer intransitive.

(92) # John is healthier than Ruth with respect to blood pressure.
Ruth is healthier than Sam with respect to blood pressure.
Sam is healthier than John with respect to blood pressure.

Indeed, comparisons involving multidimensional adjectives point out the importance of dimensions in generating vagueness effects. In examples like (93), the second premise of the sorites paradox involves comparison (*one unit higher*) along the dimension of blood pressure.

(93) Premise 1: A person whose systolic blood pressure is 100 is healthy.
Premise 2: A person whose systolic blood pressure is one unit higher than a healthy person's is healthy.
Conclusion: A person whose systolic blood pressure is 180 is healthy.

The sorites paradox thus depends on (i) the presence of a measure along some scale in the first premise and the conclusion, and (ii) a comparison (iii) along this scale in the second premise. Attempts to build a sorites paradox without these elements will fail or lose their paradoxical force (94). In (94a), the sorites paradox is missing the measures in (i). In (94b), the comparative relation in the sorites premise is replaced with an indifference relation, violating (ii). And in (94c), the sorites paradox is missing the relevant scale from (iii).

(94) a. Premise 1: A person who is very healthy is healthy.

Premise 2: A person whose systolic blood pressure is one unit higher than a healthy person's is healthy.

\# Conclusion: A person who is not healthy is healthy.

b. Premise 1: A person whose systolic blood pressure is 100 is healthy.

Premise 2: A person whose systolic blood pressure is indifferent to a healthy person's is healthy.

\# Conclusion: A person whose systolic blood pressure is 180 is healthy.

c. Premise 1: A person whose systolic blood pressure is 100 is healthy.

Premise 2: A person who is slightly less healthy than a healthy person is healthy.

\# Conclusion: A person whose systolic blood pressure is 180 is healthy.

As discussed below, the relationship between a multidimensional adjective like *healthy* and its component dimensions is crucial for deriving vagueness effects.

3.3.1 The status of dimensions in multidimensional adjectives

A "dimension" may be an explicit and quantifiable semantic argument, as in *healthy in every respect*. However, dimensions appear in a number of constructions that lack this option (*mathematically, Elena is clever*). More importantly, it is impossible to fix a scale for multidimensional adjectives without reference to some procedure for integrating many dimensions into one. Where Kennedy & McNally (2005) associate gradable adjectives with a set of degrees D, an ordering, and a dimension δ providing the property to be measured, I associate multidimensional adjectives with a set of degrees, an ordering, and a function that creates one scalar property from many, $\langle D, <, f(\delta_1, \ldots, \delta_n)\rangle$. The nature of this function is discussed in 3.4.

Kamp (1975) began the discussion of multidimensional adjectives as a challenge to supervaluationism's ability to accommodate comparatives. In particular, Kamp relied on multidimensional adjectives like *clever* to demonstrate that supervaluationism (and in particular that version proposed by Lewis (1970)) fails to capture "penumbral connections" between the positive form of an adjective (*clever*) and its comparative (*cleverer*). Kamp notes that a comparative relation defined by (95) will founder on the shoals of multidimensional adjectives like *clever.*

(95) x is at least as *Adj.* as y iff every sharpening that makes y *Adj.* makes x *Adj.*

Kamp (1975) asks: what happens when Smith, though less quick-witted than Jones, is much better at solving mathematical problems? Is Smith cleverer than Jones? "This is perhaps not clear, for we usually regard quick-wittedness and problem-solving facility as indications of cleverness, without a canon for weighing these criteria against each other when they suggest different answers" (Kamp, 1975, 140). Indeed, Kamp concludes that adjectives like *clever* are far more common than "one dimensional" adjectives like *heavy,* and "[t]here is no fixed procedure for integrating the various criteria" (p. 140).

In contrast, Sassoon (2013a) claims that the integration procedure is boolean. On the basis of the evidence in (82-85), Sassoon (2013a) treats the dimensions of multidimensional adjectives as semantic arguments. The positive use of multidimensional adjectives (96) involves implicit saturation of the dimensional argument, which the addressee presupposes in context.

(96) Ruth is healthy.

Sassoon (2013a, 338) identifies a dimension-binding operator OP that binds the different dimensions involved in the semantics of multidimensional adjectives, "creating from them one property of individuals", as in (97).

(97) $\lambda x.x$ is healthy with respect to OP dimensions

Sassoon considers candidates for OP that include what she calls "logical" operations, like quantifiers of various kinds (*all, some, most, many, none*), as well as "non-logical" operations, like *on (possibly weighted) average.*

Based on corpus studies, Sassoon (2013a) concludes that OP is classed with the logical operations. She identifies a distinction between "conjunctive," "disjunctive," and "mixed" multidimensional adjectives. Conjunctive adjectives (like *healthy*) require entities to reach the standard in *all* of their dimensions. Disjunctive adjectives (like *sick*) require entities to reach the standard in *at least one* of their dimensions. Mixed adjectives (like *intelligent*) depend on context to determine whether the relevant entity must reach the standard in all dimensions or only one dimension.

Sassoon's corpus study tracks co-occurrences of multidimensional adjectives with exceptive phrases, like *healthy except with respect to blood pressure,* because such phrases are well-known to diagnose the presence of universal quantification (as with *every* and *no*).[10]

(98) a. Everyone is happy except for Ruth.

b. No one is happy except for Ruth.

c. # Someone is happy except for Ruth.

Thus, if an adjective co-occurs with exceptive phrases, a universally quantifying operation binds its relevant dimensions—it is conjunctive. If an existentially quantifying operation binds the adjective's relevant dimensions, the adjective will not co-occur with exceptive phrases—it is disjunctive. This leads to predictions like (99) and (100) (treating a negated existential quantifier as equivalent to a universal quantifier followed by negation).[11]

10. Sassoon (2013a) relies on Hoeksema's (1995) treatment of exceptive phrases, which requires that the sentential argument of an exceptive phrase (the *everyone is happy* portion of (98a), for instance) be equivalent to a universally quantified statement.

11. Sassoon draws support for her interpretation of OP from other sources as well, including antonymy and scale type (e.g., partially closed vs. totally closed, see 2). For example, she hypothesizes that an adjective is conjunctive if and only if its antonym is disjunctive, as with *healthy* and its antonym *sick*. She also hypothesizes that conjunctive adjectives will associate with upper-closed scales, that disjunctive adjectives will associate with lower-closed scales, and that mixed adjectives will associate with open scales, as illustrated by the diagnostics for scale type in (1).

(99) a. Stephen is healthy except for high blood pressure.

b. # Stephen is sick, except for high blood pressure.

(100) a. John is not sick, except for the flu.

b. # John is not healthy, except for the flu.

Sassoon's study establishes potential default interpretations for conjunctive and disjunctive multidimensional adjectives (mixed multidimensional adjectives do not have such default interpretations).

However, the study does not settle whether the dimension-binding operator OP—the procedure, in Kamp's (1975) terms, for integrating the relevant criteria—is "logical" or something else. First, supposedly conjunctive multidimensional adjectives like *healthy* are compatible with non-logical dimension-binding operators (using Sassoon's terminology), as well as "disjunctive" dimension-binding operators:

(101) a. Clarence is healthy on average.

b. Clarence is healthy in one respect.

This suggests that exceptive phrases may induce—rather than diagnose—the existence of a universal quantifier over dimensions for a conjunctive adjective like *healthy*.

Second, the exceptive-phrase test produces less reliable results with multidimensional adjectives than it does with DP quantifiers like *every*. Sassoon's test for conjunctive vs. disjunctive dimension-binding procedures ultimately boils down to whether a multidimensional adjective with an exceptive phrase is used in "positive" or "negative" context, as defined in the following way ($\textsc{dim}(P)$ represents the dimensions associated with the predicate P):

---

(1) a. entirely healthy / #entirely sick

b. #slightly healthy / slightly sick

The latter source of evidence is suggestive, but ultimately Sassoon, 374–5 concludes that the data do not support strong generalizations based on scale type. In later work (Sassoon, 2015), Sassoon does not build the relevant quantifiers into the denotations of multidimensional adjectives.

(102) a. A context C is positive iff hits of the form "C P except D'" are roughly equivalent to $\lambda x.\forall Q \in \textsc{dim}(P) - P$ with respect to D: x is Q and x is not P with respect to D.

b. A context C is negative iff hits of the form "C P except D'" are roughly equivalent to $\lambda x.\forall Q \in \textsc{dim}(P) - P$ with respect to D: x is not Q and x is P with respect to D.

Conjunctive adjectives are used more often in positive contexts than in negative contexts, and disjunctive adjectives are used more often in negative contexts than in positive contexts. Yet while this finding confirms an overall tendency, Sassoon, 379-380 also shows that it is not categorical, or at least not as categorical as other judgments involving exceptive phrases and quantifiers like *every*. For example, the infelicity of (103) is stark, but at least in some cases, *healthy* felicitously co-occurs with an exceptive phrase in a negative context (as in Sassoon's (104)).

(103) # Not everyone is happy, except for Ruth

(104) It should be noted that many affluent neighborhoods, also, are not healthy, except for their economic condition.

Third, the category of "mixed" multidimensional adjectives (like *intelligent*) itself suggests a role for non-logical dimension-binding operations in the semantics of multidimensional adjectives.

Compositionally, Sassoon (2015) illustrates how a dimension argument might affect the semantics of comparatives. Sassoon describes a sentence like (105) as ambiguous in two ways.

(105) Nino is healthier than Tony.

On one interpretation, Nino exceeds Tony with respect to some degree on a dimension of health, like cholesterol. But on another reading, the number of dimensions in which Nino is healthy exceeds the number of dimensions in which Tony is healthy. This second reading Sassoon calls the "dimension counting" reading.

To account for this reading, Sassoon introduces a dimension argument, which can be bound by an explicit quantifier like *every, some,* etc. Sassoon takes gradable adjectives to be relations between entities and degrees, following von Stechow (1984). The positive form of a gradable adjective combines with a phonologically silent degree modifier *pos*, which provides the adjectival standard. Gradable adjectives themselves are relations between entities (type e) and degrees (type d). D represents a set of degrees (type $\langle d, t\rangle$).

(106) a. $[\![pos]\!] = \lambda D_c \lambda D.\forall d \in D_c[D(d)]$ (the "membership norm" D_c is an interval)

b. $[\![\text{tall}]\!] = \lambda x \lambda d.\text{tall}(x, d) = x$ is tall to degree d

c. $[\![\text{I}_c]\!]$ = an interval that functions as the "contextual standard interval" of an adjective

d. $[\![$Ruth is *pos* tall$]\!]$

$= \text{pos}(I_c)(\lambda d.\text{tall}(\mathbf{r}, d))$

$= [\lambda D.\forall d \in I_{\text{tall}}(D(d))](\lambda d.\text{tall}(\mathbf{r}, d))$

$= \forall d \in I_{\text{tall}}[\text{tall}(\mathbf{r}, d)]$

"For every degree d in *tall*'s standard interval, Ruth is tall to at least degree d."

Multidimensional adjectives are ambiguous between the representation in (107a) and the representation in (107b) (the "dimension-counting" denotation for *healthy*).

(107) a. $[\![\text{healthy}]\!] = \lambda x \lambda d.\text{healthy}(x, d)$

b. $[\![\text{healthy}]\!] = \lambda n \lambda x.|\lambda R \in \textsc{dim}_{\text{healthy}}.\forall d \in I_R(R(x, d))| \geq n$

In (107b), $\textsc{dim}_{\text{healthy}}$ is the set of degree relations denoted by the dimensions of the adjective *healthy*, and n is a "degree" on a scale of cardinality of subsets of $\textsc{dim}_{\text{healthy}}$. R is a predicate of type $\langle d\langle e, t\rangle\rangle$, the type of gradable adjectives. The idea is that *Stephen is healthy* is true if Stephen is healthy on some contextually determined number of health dimensions.

(108) $[\![\text{Stephen is } pos \text{ healthy}]\!]$

$= \text{pos}(I_{\text{healthy}})(\lambda n.|\lambda R \in \text{DIM}_{\text{healthy}}.\forall d \in I_R(R(\mathbf{s}, d))| \geq n)$

$= [\lambda D.\forall d \in I_{\text{healthy}}(D(d))](\lambda n.|\lambda R \in \text{DIM}_{\text{healthy}}.\forall d \in I_R(R(\mathbf{s}, d))| \geq n)$

$= \forall d \in I_{\text{healthy}}|\lambda R \in \text{DIM}_{\text{healthy}}.\forall d' \in I_R(R(\mathbf{s}, d'))| \geq d$

"Stephen is healthy on at least d dimensions, where d is the number of dimensions of health in *healthy*'s standard interval."

Now Sassoon is in a position to capture the ambiguity of sentences like *Bill is healthier than Ann in every respect.* A PP like *with respect to* must adjoin to the multidimensional adjective in order to introduce the dimension argument into the derivation. The PP denotes a modifier as in (109). It takes two arguments, R the relevant dimension (like cholesterol of flu for *healthy*) and the dimension-counting denotation (107b) of a multidimensional adjective R_A. The modifier presupposes that R is a dimension of the relevant multidimensional adjective. Semantically, R is of the same type as a gradable adjective (a relation between entities and degrees). Sassoon's derivation of *healthy with respect to the flu* is in (110).

(109) $[\![\text{with respect to / in / in terms of}]\!] = \lambda R \lambda R_A : R \in \text{DIM}_A.R$

(110)

healthy with respect to the flu

$\lambda d \lambda x.R_{\text{flu}} \in \text{DIM}_{healthy}.R_{\text{flu}}(x, d)$

healthy

$\lambda n \lambda x.|\lambda R \in \text{DIM}_{\text{healthy}}.\forall d \in I_R(R(x, d))| \geq n$

$\lambda R_A : R_{\text{flu}} \in \text{DIM}_A.R_{\text{flu}}$

with respect to

$\lambda R \lambda R_A : R \in \text{DIM}_A.R$

flu

R_{flu}

Finally, Sassoon defines quantification over dimensions, as in expressions like *healthy in every respect.* In order to resolve a type mismatch, the quantified expression *every respect*, $\lambda Q_{\langle\langle d\langle e,t\rangle\rangle t\rangle}.\forall R[Q(R)]$, raises (via quantifier raising), leaving a trace of type $\langle d\langle e, t\rangle\rangle$. The derivation proceeds, per Sassoon, as in Fig. 3.3.1.

Figure 3.4: John is healthy in every respect

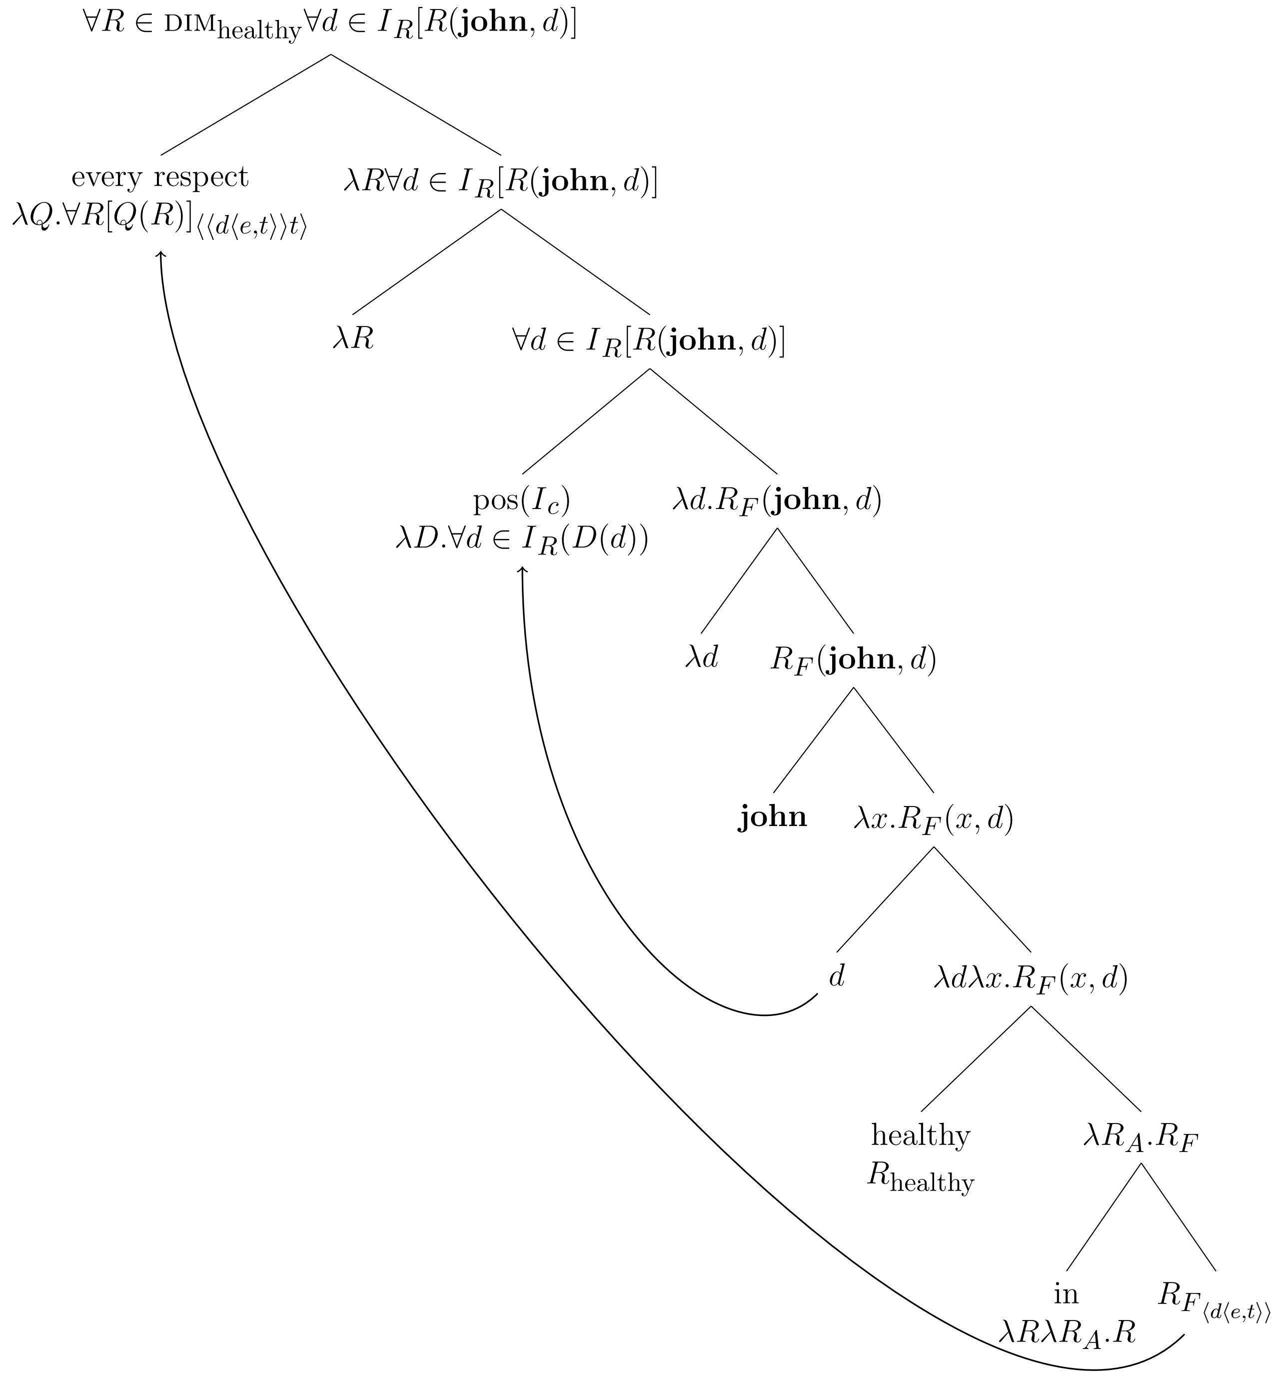

On Sassoon's (2015) picture, then, dimensions are not integrated; rather an entity is evaluated along each dimension separately.[12] This means that Sassoon cannot explain how the scale of healthiness comes into being for (107a), repeated below, or related comparatives.

(107a) $[\![\text{healthy}]\!] = \lambda x \lambda d.\text{healthy}(x, d)$

Assuming a default "dimension-counting" for *healthy*, as in (107b), also poses problems.

(107b) $[\![\text{healthy}]\!] = \lambda n \lambda x.|\lambda R \in \textsc{dim}_{\text{healthy}}.\forall d \in I_R(R(x, d))| \geq n$

First, as Bale (2008) argues, comparisons of multidimensional adjectives are not generally exercises in dimension-counting. Instead, Bale proposes a universal scale of measurement, such that multidimensional adjectives like *intelligent* are functions from individuals to degrees on this universal scale ("universal degrees"). This allows for comparatives like (111), which compare relative positions on separate scales of intelligence and beauty.

(111) Sonia is more intelligent than Sam is beautiful.

Second, it is unlikely that "respects" function like gradable adjectives. That is, (112a) may act like a gradable adjective, but (112b) almost certainly does not.[13]

(112) a. healthy with respect to the flu

b. (with respect to) the flu

Third, respect-accessing operations may be introduced in a variety of ways, not only with PPs. While Sassoon (2015) does not expressly limit how respects may be introduced into the derivation, the treatment of respects as gradable adjectives creates challenges for an account of the data in (113b) and (113c).

(113) a. Elena is clever with respect to mathematics.

b. Mathematically, Elena is clever.

12. See Sassoon (2015) for an explanation of the denotation of *in* in Fig. 3.3.1.

13. Compare Sassoon (2015, 19): "The NP complement of a respect-accessing preposition denotes a respect (R_{flu}) of the multidimensional adjective to which the PP projective adjoins."

c. Elena is clever, math-wise.

Not all of the equivalent meanings (113a-113c) need arise the same way. However, they all appear to involve respect-accessing operations, and this means that respects—whatever they are—enter into the derivation in contexts without PPs. Sassoon's (2015) treatment is not readily generalizable to apparently related cases.

The related meanings in (113a-113c) suggest that respects—what I have called dimensions—are accessible in a variety of contexts and in a variety of ways. Included in this list are (114a-114c), constructions that include comparatives (114a), topics (114b), and modals.

(114) a. Elena is cleverer in math than Sam is.

b. (Elena is a musical dunce, but) as for math, Elena is clever.

c. (In view of what the law requires,) Elena must be reasonable.

Given the broad accessibility of dimensions across constructions, I take dimensions to be part of a multidimensional adjective's denotation. Sometimes dimensions are introduced as explicit semantic arguments (113a), but sometimes not (113b, 114b). Accordingly, multidimensional adjectives are not ambiguous between (107a) and (107b)—they have the same denotation as other gradable adjectives (115). And as Kamp (1975) noted, multidimensional adjectives involve an "integration procedure" to get from multiple dimensions to a single scale. I will refer to this integration procedure as an "aggregation function," for reasons that become clear below.

(115) $\lambda x \lambda d.\text{healthy}(x, d)$

3.4 Multidimensional adjectives as collective choice functions

The function f in the multidimensional adjectival scale $\langle D, <, f(\delta_1, \ldots, \delta_n)\rangle$ is a collective choice function: it aggregates many rankings (or, equivalently, choices) into one. For instance, *healthy* is a measure function that maps an entity to a position on a one-dimensional scale of healthiness, as in Fig. 3.2 (replacing "height" with "healthiness" and *tall* with *healthy*). But how is the one-dimensional healthiness scale constructed from the multiple dimensions of healthiness? As Kamp (1975) suggested, measures along cholesterol, blood pressure, or other contextually relevant dimensions are somehow integrated to produce a single property measured by the relevant scale.[14]

Moreover, we have intuitions about how dimensional measures should affect the meaning of the multidimensional adjective. For instance, if every dimension involved in a use of *healthy* ranks Elena above others, then a proper use of *healthy* should reflect this ranking. Sentences like (116) should therefore seem semantically odd.

(116) # Elena is healthy in every respect, but she is not healthy.

The oddness of (116) reflects Arrow's principle of Unanimity: if every dimension ranks Elena strictly higher than other members of the comparison class, the aggregated ranking should also rank Elena highly.

Similarly, it should be possible to determine the relative health of Nino and Ruth by considering only the relative rankings of Nino and Ruth along the relevant dimensions. A phenomenon known as "implicit comparison" may verify this effect. Imagine a dialogue as in (117).

(117) Speaker A: Compared to Ruth, Nino is healthy.
Speaker B: But consider Clarence.
Speaker A: # Then, compared to Ruth, Nino is not healthy.

If Speaker A's conclusion in (117) is anomalous, then this is a linguistic analog of Independence of Irrelevant Alternatives. In (117), Clarence is an irrelevant alternative.

We also have no reason to believe that some possible rankings of members of the comparison class should be excluded at the outset. I take this to support the initial plausibility of Unrestricted Domain.

14. The single-scale also requirement is important for deriving the semantics of comparatives like *Elena is healthier than Ruth*.

An analog of Non-Dictatorship may be inferred from the oddness of (118).

(118) (Scenario: Sonia has good blood pressure, but terrible cholesterol and heart rate.)

Sonia is healthy.

In (118), the speaker's conclusion is odd because one dimension does not usually dictate the meaning of a multidimensional adjective. If the dimension of blood pressure were "dictatorial" in the context represented by (118), the speaker's conclusion should strike us as natural. Similarly, the speaker's conclusion in (119) may be natural, despite the middling ranking of Sonia with respect to blood pressure.

(119) (Scenario: Sonia has middling blood pressure, but good cholesterol and heart rate)

Sonia is healthy.

Taken together, (118) and (119) suggest that one dimension does not usually control the interpretation of a multidimensional adjective in a "dictatorial" fashion.

If gradable adjectives are choice functions, then they are defined with respect to a choice rule.

(120) $[\![\text{healthy}]\!]^c = \{x \mid \forall y \in c,\ xRy\}$

For multidimensional adjectives, this choice rule is a function of the choice rules used by the multidimensional adjective's component dimensions.

(121) $R_{\text{healthy}} = f(R_{\text{bp}}, R_{\text{cholesterol}}, R_{\text{heart rate}}, \ldots)$

Arrow (1950) and Chichilnisky (1982) tell us that, under certain assumptions, the mapping f may result in a choice rule that is intransitive and discontinuous.

For example, it is easy to recreate Condorcet's paradox in the adjectival context, replacing voters with dimensions and candidates with members of the comparison class. Suppose there are three members of the comparison class (John, Steven, and Ruth) and three dimensions (blood pressure, cholesterol, and heart rate) contextually relevant to the interpretation of *healthy*, as in (122).[15]

15. The application of Arrow's Theorem to individual decision making has a lengthy pedigree, including

(122)

bp:	john	>	ruth	>	stephen
cholesterol:	stephen	>	john	>	ruth
heart rate:	ruth	>	stephen	>	john

This table should look familiar. As Égré & Klinedinst (2011, 10) rightly worry, multidimensional adjectives may possibly give rise to "intransitivities of a kind familiar in social choice theory (a may be judged more clever than b along one dimension, b more clever than c along another, and c more clever than a along a third)."

Indeed, it is possible to construct Arrow's paradox within adjectival semantics by slightly reframing Arrow's axioms. Within an interpretive context, fix a finite set $X = a, b, c, \ldots, n$ of members of the comparison class and a finite set $D = 1, 2, 3, \ldots, i$ of relevant dimensions. Let $\succcurlyeq_i$ be the "dimension ordering" of dimension i with respect to members of the comparison class. Thus $a \succcurlyeq_i b$ means "dimension i orders alternative a at least as high as alternative b." A dimension ordering is transitive and complete.

Let the relation $\succcurlyeq_{\text{Adj}}$ represent the "adjectival ordering" derived from aggregating the set of dimension orderings $\succcurlyeq_1, \succcurlyeq_2, \ldots, \succcurlyeq_i$. That is, $\succcurlyeq\, = f(\succcurlyeq_1, \succcurlyeq_2, \ldots, \succcurlyeq_i)$. For example, the relation $\succcurlyeq_{\text{healthy}}$ represents the ranking of entities with respect to healthiness based on aggregating measures along the dimensions $\succcurlyeq_{\text{bp}}, \succcurlyeq_{\text{heart rate}}, \ldots, \succcurlyeq_{\text{cholesterol}}$. The entities considered *healthy* are those at the top of the ranking.[16]

The Arrowian axioms reformulated for adjectival semantics are presented in formal terms below.

(123) a. Unanimity: For all i and every pair of alternatives a, b in X, if $a \succcurlyeq_i b$ for all dimensions i, then $a \succcurlyeq_{\text{Adj}} b$.

b. Independence of Irrelevant Alternatives: For every pair of entities a and b and every pair of "profiles" $(\succcurlyeq_1, \succcurlyeq_2, \ldots, \succcurlyeq_i), (\succcurlyeq_{1'}, \succcurlyeq_{2'}, \ldots, \succcurlyeq_{i'})$, if for all dimensions i the

work by Arrow himself. See Arrow & Raynaud (1986); Greene et al. (2001); Livnat & Pippenger (2006); Weyl (2009); Okasha (2011); Katz (2011); Stegenga (2013); Morreau (2015).

16. Whether this is the top of the ranking or "near" the top will not matter in light of Arrow's result.

relations $\succsim_i$ and $\succsim_{i'}$ coincide on a, b, then $\succsim_{\text{Adj}}$ and $\succsim_{\text{Adj}'}$ coincide on a, b. By "coincide" I mean that both relations rank a over b, both rank b over a, or both rank a and b equally.

c. Unrestricted Domain: The domain of the aggregation function f potentially includes every list $\langle \succsim_1, \succsim_2, \ldots, \succsim_i \rangle$ of i weak orderings of the alternatives in X.

d. Non-Dictatorship: There is no dimension i such that $f(\succsim_1, \succsim_2, \ldots, \succsim_i) = \succsim_i$ for all profiles $(\succsim_1, \succsim_2, \ldots, \succsim_i)$.

Multidimensional adjectives appear to follow analogues to the Arrowian assumptions Unrestricted domain, Unanimity, Independence of Irrelevant Alternatives, and Non-dictatorship/anonymity. Therefore, following a similar strategy that was used in section 3.1, it is possible to prove that choice rules associated with multidimensional adjectives like $\succsim_{\text{healthy}}$ may be intransitive and discontinuous.

3.5 Vagueness effects following from limitations on collective choice

Arrow and Chichilnisky's results show that seemingly sensible constraints on f—like Unanimity, Unrestricted Domain, Independence of Irrelevant Alternatives, and Non-Dictatorship—may lead f to fail in particular ways. Either the constraints themselves will be violated or other desirable properties like transitivity will be violated. Vagueness effects like the sorites paradox and borderline cases flow from these failures.

First, the role of intransitivity in vagueness effects is well-explored by Cobreros et al. (2012), who claim that "the non-transitivity of the indifference relation is a central feature of all vague predicates" (p. 349). Cobreros et al. (2012) build a theory of vagueness upon such intransitivities, training their analysis on the second premise of the sorites paradox (124), which represents the idea that vague predicates are tolerant (125).

(124) Premise 1: A person whose systolic blood pressure is 100 is healthy.

Premise 2: A person whose systolic blood pressure is one unit higher than a healthy person's is healthy.

Conclusion: A person whose systolic blood pressure is 180 is healthy.

(125) Premise 2: $\text{healthy}(a) \wedge a \sim_{\text{healthy}} b \rightarrow \text{healthy}(b)$

More generally, a tolerant predicate P obeys (126) for "indistinguishable" entities x_i and x_{i+1}.

(126) $\forall x_i, x_{i+1} : (P(x_i) \wedge x_i \sim^{P} x_{i+1}) \rightarrow P(x_{i+1})$

For example, the relation "indifferent with respect to health" is reflexive and symmetric, but not transitive: Elena's health maybe observationally indistinguishable from Sonia's (to use Kamp's (1975) term), and Sonia's may be observationally indistinguishable from Nino's, but Elena and Nino may have distinct measures of healthiness.

Tolerance implies both the transitivity of the indifference relation and the idea that small changes along one dimension of a multidimensional predicate result in small changes in the denotation of the multidimensional predicate. First, tolerance implies transitive indifference. For example, (126) implies that if $x_i \sim x_{i+1}$ and if $x_{i+1} \sim x_{i+2}$, then $x_i \sim x_{i+2}$ for all x_i. If *healthy* is tolerant—as the second premise of (124) appears to suggest—then the indifference relation associated with the choice rule R_{healthy} is transitive. However, R_{healthy} is the product of a collective choice function f (121), and limitations on collective choice entail that f might not result in a transitive collective choice. The indifference relation $\sim_{\text{healthy}}$ may therefore be intransitive. Since tolerance implies the transitivity of the indifference relation, and since collective choice shows that this relation might not be transitive, tolerance is false.

Indeed, it is possible to trace intransitivity to a particular Arrowian axiom. Independence of Irrelevant Alternatives restricts the collective choice function f in ways that give rise to intransitivity. As Saari (2008, 44) notes, if the aggregation function f obeys Independence, f may be rewritten as (127), where each component of f is determined by f's ranking of a pair of entities.

(127) $f = (f_{\{\text{Elena, Sonia}\}}, f_{\{\text{Sonia, Nino}\}}, f_{\{\text{Nino, Elena}\}})$

It is easy to see that such a rule allows for intransitive outcomes: f in (127) can rank Elena indifferent to Sonia, Sonia indifference to Nino, and Nino over Elena, without contradiction.

Other theories of vagueness incorporate something very much like Independence. Fara (2002, 59) relies on the Similarity Constraint in (128) to derive vagueness effects.

(128) **Similarity Constraint.** Whatever standard is in use for a vague expression, anything that is saliently similar, in the relevant respect, to something that meets the standard itself meets the standards; anything saliently similar to something that fails to meet the standard itself fails to meet the standard.

The way the Similarity Constraint works in practice, as Fara explains with respect to the vague predicate *tall*, is the following: "for any particular x and y that differ in height by just 1mm, the very act of our evaluation raises the similarity of the pair to salience, which has the effect of rendering *true* the very instance we are considering." Thus, Fara affirms each instance of the sorites premise (*if* n *is* P, *then* $n+1$ *is* P) without affirming the universally-quantified version ($\forall n$ *if* n *is* P, *then* $n+1$ *is* P). The sorites premise is therefore false.

Fara's (2002) Similarity Constraint likewise embodies the principle of Independence. This constraint requires "salient" comparisons, which turns out to mean pairwise comparisons. Fara rules out the universally quantified sorites premise because each comparison cannot be simultaneously salient to the speaker: "there are too many pairs for us to actively entertain each similarity" (p. 69).

Second, tolerance implies continuity.[17] If (126) is true, then small changes along one

17. The definition of continuity is dependent on the particular topology imposed on the choice space (that is, what constitutes the "open" sets).

(1) **Continuous map**: Let X and Y be topological spaces. A map $f : X \to Y$ is called continuous if the inverse image of opens sets is always open. (Janich, 1984)

"[A]ny reasonable topology, certainly any topology which has been used on preference spaces" reduces to a Euclidean choice space (Heal, 1997, 3), in which continuity may be defined as follows:

dimension of health, like blood pressure, result in small changes in the value of *healthy.* In particular, where x_i and x_{i+1} differ only by one unit of systolic blood pressure, either both are healthy ($[\![\text{healthy}(x_i)]\!]=1=[\![\text{healthy}(x_{i+1})]\!]$) or both are not ($[\![\text{healthy}(x_i)]\!]=0=[\![\text{healthy}(x_{i+1})]\!]$). However, R_{healthy} is the product of a collective choice function f (121), and limitations on collective choice entail that f might not result in a continuous collective choice. Since tolerance implies a continuous relationship between dimensional measures and the meaning of a multidimensional predicate, and since collective choice shows that this relation might not be continuous, tolerance is false.

Importantly, Smith (2008) notes that continuity is too strong a constraint to model tolerance. Continuity overgenerates the number of vague predicates in part because the definition of continuity is so flexible.[18] Topologically speaking, every discrete domain could support a continuous function, but this elides important distinctions between vague predicates like *healthy* and non-vague predicates like *has 100 hairs or fewer.* Thus, Smith adopts a "local" notion of continuity that is equivalent to the intuitive description of continuity I have been

---

(2) **Continuity:** for any map $f : \mathbb{R}^m \to \mathbb{R}^k$, the map f is continuous at $p \in \mathbb{R}^m$ if for each $\epsilon > 0$, there exists a $\delta > 0$ such that $|f(q) - f(p)| < \epsilon$ as soon as $|q - p| < \delta$.

18. In order to define continuity, we need to impose a topology on the set constituting the domain of the predicate P. For example, if a subset S of the domain satisfies the following property, S is a basis element of the relevant topology.

(1) $(x \leq {}_z y \wedge y \in S \wedge z \in S) \to x \in S$

That is, if x is at least as close to z as y is, and y and z are both in S, then x is in S. (This is Smith's three-place similarity relationship, but he makes a convincing argument that a two-place similarity relationship works similarly.) Then $\forall x, y(x \leq {}_y x)$, which means that every singleton set is a basis element of the topology, and, consequently, every function between sets of entities is continuous. Such a topology fails to discriminate between linguistically vague predicates and predicates like *has 100 hairs or fewer.* The latter is not vague. Yet if we impose the topology generated by (1) on the domains of all predicates, all predicates will turn out to be vague, an unwelcome result.

using: a function is locally continuous if a *small* change in input produces at most a *small* change in the value of the function.

Instead, Smith argues for a "local" notion of continuity, which he dubs "closeness." Vague predicates support local continuity (a *small* change in input produces at most a *small* change in the value of the function) but *not* "global" continuity, which Smith (2008: 155) describes as the following: "for *any* positive-sized target area in the co-domain (whether or not we would ordinarily regard it as 'small') we can find a positive-sized launch area in the domain (which again need not be 'small' in the ordinary absolute sense) such that *everything* sent by the function from that launch area lands in that target area."

Similarly, Weber & Colyvan (2010) argue for an understanding of tolerance as something like local continuity. According to them, "A predicate is vague iff its characteristic function is locally constant but not globally constant" (p. 318), where local constancy is understood in a way similar to "local" continuity (129). Indeed, the authors write, "[s]ubstituting 'continuous' for 'constant' . . . would make no great difference," (fn. 17) acknowledging Smith's objections to global continuity.

(129) A function f is locally constant iff for each $x \in X$ there is a neighborhood of x U_x such that the restriction of f to U_x is constant.

Weber & Colyvan (2010: 325) explain that "the core notion of local constancy is a generalization of the principle of tolerance." And Rizza (2013), critiquing this approach, suggests that "sorites-type arguments [are] based on a dichotomy between a local and global level."

I agree with Rizza's diagnosis, at least as understood through the lens of social choice. Chichilnisky (1982) shows that continuous dimensions (the local level) do not necessarily map into continuous predicates (the global level).[19] And this difficulty, in turn, can be traced to the nature of multidimensional choice space.

19. See chapter 4 for a more extended discussion of Chichilnisky's (1982) result.

3.6 The curse of dimensionality

Saari (2008, 13) refers to the pathologies of collective choice in a multidimensional choice space as "the curse of dimensionality." Collective choice functions lose information when mapping multiple "local" dimensional rankings (like the ranking of entities according to their blood pressure) into a "global" aggregated ranking (like the scale of healthiness). This lost information ultimately results in vagueness effects like those described above.

In one sense, independence plays an information-losing role in the semantics of gradable adjectives. By hypothesis, each component dimension of a multidimensional adjective (like blood pressure for *healthy*) is a linear order: transitive and complete. If the aggregation function f of a multidimensional adjective obeys Independence, then the aggregation function is equivalent to one that considers only pairs of entities. But then the aggregation function f ignores the fact that the component dimensions are transitive. This lost information leads to an intransitive indifference relation associated with the multidimensional adjective, and an intransitive indifference relation is associated with vagueness effects. Similarly, Saari (2008) argues that continuity it itself a local concept that "ignores" global structure, and "this feature of ignoring the global structure is the same local versus global problem that arises with" Independence (p. 57).

With a little squinting, it is clear that this dichotomy pervades other accounts of vagueness phenomena, as well. In Cobreros et al.'s (2012) many-valued approach, for example, there are three notions of truth: classical truth, tolerant truth, and strict truth. Tolerant truth is a weaker notion of truth than classical truth, and strict truth is a stronger notion. The local inferences—the inferences that make up the second premise of the sorites argument—are tolerantly true. The idea is that you can only string so many tolerantly true inferences together (at the same time) before you end up with a false conclusion.[20] For example,

20. Other many-valued logics share this feature. For example, infinitely-valued logics often use a value along the interval $[0, 1]$ as the "degree of truth" of a certain proposition (Zadeh, 1975; Hyde, 2008). Infinitely-valued logics are often defined in such a way that an argument is only valid if it preserves the maximum

Cobreros et al. (2012, 376) validate each individual step in the sorites premise (130a), but not the universally-quantified version (130b).[21]

(130) a. Premise 1: $P(a)$

Premise 2: $a \sim b$

Conclusion: $P(b)$

b. Premise 1: $P(a)$

Premise 2: $\forall i \in [i, n-1]\ a_i \sim a_{i+1}$

Conclusion: $P(n)$

Similarly, contextualists like Fara (2002) affirm each instance of the sorites premise (*if n is P, then* $n+1$ *is P*) without affirming the universally-quantified version ($\forall n$ *if n is P, then* $n+1$ *is P*). As Smith (2008, 113) describes contextualist accounts, "The basic idea behind contextualism is that vagueness is a diachronic phenomenon, which only emerges when we consider the semantic state of a language over time (or more generally, over multiple instances of interpretation)." Contextualist accounts place particular emphasis on the act of interpretation, arguing that the act itself changes the semantic "facts on the ground." For example, Fara (2002, 59) explains that "for any particular x and y that differ in height by just 1mm, the very act of our evaluation raises the similarity of the pair to salience, which has the effect of rendering *true* the very instance we are considering." Like other contextualist accounts (Kamp, 1981; Raffman, 1994; Shapiro, 2006), Fara affirms each instance of the sorites premise (*if n is P, then* $n+1$ *is P*) without affirming the universally-quantified

degree of truth in the argument's premises. The sorites argument fails to be valid on this account because each step in the second premise, while "almost true" itself, reduces the degree of truth of the conclusion bit by bit. In this way, the first premise is perfectly true, each instantiation of the tolerance principle is *almost* perfectly true, and the conclusion is perfectly false. Since the sorites argument fails to preserve the maximum degree of truth, it is invalid.

21. The key to the validity of (130a) and the simultaneous invalidity of (130b) is an intransitive notion of logical consequence.

version ($\forall n$ *if* n *is* P*, then* $n+1$ *is* P).

Tellingly, if we consider each instance of the sorites premise simultaneously, the sorites paradox resurfaces. Fara's response is that such simultaneous consideration is impossible. Each comparison cannot be simultaneously *salient* to the speaker: "there are too many pairs for us to actively entertain each similarity" (p. 69).

As Gaertner (2009) demonstrates, however, there is an additional interpretation of Arrow's theorem involving information loss. Arrow's approach depends on *ordinal* dimensional rankings in a multidimensional choice space, and this imports certain restrictions on the information available to the aggregation function.

An *ordinal scale* is a structure $(S, \succsim)$ where $\succsim$ is a weak order. Ordinal scales are too informationally impoverished to allow numerical measurements like *two meters* or ratio measurements like *twice as*. Moltmann (2005) uses this as a basis for proposing an ordinal analysis for multidimensional adjectives, since adjectives like *clever* and *happy* are also incompatible with numerical and ratio modifiers (Sassoon, 2013b).[22]

An ordinal scale can be mapped to a numerical representation in any way that respects the weak order $\succsim$. For example, take a structure preserving function μ from S into the real numbers $\mathcal{R}$. The following numerical representations are equivalent, so far as the ordinal scale is concerned.

(131) μ is a homomorphism from $(S, \succsim)$ to $\langle \mathbb{R}, \geq \rangle$ such that $a \succsim b$ iff $\mu(a) \geq \mu(b)$

$a \succsim b$ iff $\mu(a) = 10$ and $\mu(b) = 20$

$a \succsim b$ iff $\mu(a) = 0.13$ and $\mu(b) = 0.14$

$a \succsim b$ iff $\mu(a) = 100,000$ and $\mu(b) = 9,000,000$

22. Other partisans of the ordinal scale approach include Cresswell (1976a); Klein (1980). As Sassoon (2013b, 80) puts it, "The ordinal approach takes the ordering relation between entities . . . to be conceptually primitive, and the degrees . . . to be derived from them." The derivation is in terms of equivalence classes: degrees are equivalence classes under the weak order $\succsim$, where an equivalence class is the set of entities related by a reflexive, transitive, and symmetric equivalence relation (for example, a degree of *height* is an equivalence class under the equivalence relation *equally tall*).

That is, the ordered equivalence classes can be associated with any numerical representation that respects the ordering.

This property of ordinal scales leads to Arrow's result (Blackorby et al., 1984; Gaertner, 2009) in a multidimensional choice space. It is possible to turn a transitive and complete binary relation like $\succcurlyeq$ (see 60-61) into a continuous function d (132).[23]

(132) d: for each $x \in X$, $\{x \in X | x \succcurlyeq x'\}$ is closed and $\{x \in X | x' \succcurlyeq x\}$ is closed.

We can call this continuous function a "dimension function" d_i, in analogy to the dimension ordering $\succcurlyeq_i$. Each dimension evaluates a member of the comparison class $x \in X$ in terms of its dimension function $d_i(x)$.

Thus, in this revised framework, the domain of the aggregation function is sets of n-tuples of dimension functions $(d_1, d_2, \ldots d_n)$, each called a profile D. We can think of dimension functions d_i as measure functions for dimension i—each dimension measures a member of the comparison class $x \in X$ in terms of its dimension function d_i. For example, a profile relative to a use of *healthy* may contain the ordered tuple $(60, 100, 190)$ for $(d_{\text{heart rate}}(x), d_{\text{blood pressure}}(x), d_{\text{cholesterol}}(x))$ for some member of the comparison class $x \in X$. The aggregation function f (what Blackorby et al. call a "functional") is a mapping from these tuples to an ordering of X, what was called the "adjectival ordering" $\succcurlyeq_{\text{Adj}}$ above and what I'll call R_{Adj} in this context. That is, $f(D) = R_{\text{Adj}}$.

The Arrowian axioms have analogues in this revised framework. Unrestricted Domain is admitting all logically possible n-tuples. Independence of Irrelevant Alternatives has an unchanged meaning: if for any two members of the comparison class $x, y \in X$ and two profiles D and D', x and y obtain the same n-tuple of rankings in D and D' ($D(x) = D'(x)$ and $D(y) = D'(y)$), then R_{Adj} and $R_{\text{Adj}'}$ must coincide on $\{x, y\}$.[24]

Finally, we introduce a condition similar to Unanimity:

23. The analogue in the Arrowian framework is a utility function. A set is closed if it contains its boundary.

24. $D(x) = D'(x)$ means that $d_i(x) = {d'}_i(x)$ for all i.

(133) Pareto Indifference: For all $x, y \in X$ and for all profiles D, if $D(x) = D(y)$, then $xI_{\text{Adj}}y$.

This condition holds that if all rankings are indifference between two members of X, then the adjectival ordering should also be indifference (represented as $xI_{\text{Adj}}y$).

Together, Unrestricted Domain, Independence of Irrelevant Alternatives, and Pareto Indifference enforce the idea that the aggregation function f ignores all features of the relevant alternatives except how they are ranked by the relevant dimensions (Sen, 1977). This has the result, as shown by Blackorby et al. (1984), of allowing us to consider an ordering of $R*$ (of $\mathbb{R}^n$) that orders vectors of individual dimension functions directly. That is,

(134) If f satisfies Unlimited Domain, Independence of Irrelevant Alternatives, and Pareto Indifference, then there exists an ordering $R*$ of $\mathbb{R}^n$ such that $\forall x, y \in X$ and all profiles D, $xR_{\text{Adj}}y \iff \bar{d}R*\bar{\bar{d}}$, where $\bar{d} = D(x)$ and $\bar{\bar{d}} = D(y)$.

The idea is that $\bar{d}$ and $\bar{\bar{d}}$ represent vectors of measures along the relevant dimensions for every member of the comparison class x and y. $R*$ orders these vectors rather than members of the comparison class, but the orderings R_{Adj} and $R*$ are provably equivalent.

Blackorby et al. (1984) show that the advantage of this move is that whatever properties f—the aggregation function—has, $R*$ inherits.[25] The authors then show that $R*$ is a dictatorship. In present terms, whenever some individual dimension $\bar{d}_i$ has a higher measure under $\bar{d}$ than under $\bar{\bar{d}}_i$, $\bar{d}$ is ranked higher than $\bar{\bar{d}}$.

Blackorby et al. (1984) rely on a figure like Fig. 3.5 for their proof. In Fig. 3.5, two dimensions are involved, dimension 1 (d_1) and dimension 2 (d_2), and these dimensions are the components of the vector $\bar{d}$. For example, $(4, 5)$ represents a value $\bar{d}$ where $\bar{d}_1 = 4$ and $\bar{d}_2 = 5$.

In Fig. 3.5, we consider the plane's four regions. Both d_1 and d_2 rank all points in region I above $\bar{d}$, and both dimension 1 and dimension 2 rank $\bar{d}$ above all points in region III. We

25. Thus, $R*$ also behaves the axioms Unlimited Domain, Independence of Irrelevant Alternatives, and Pareto Indifference (see Blackorby et al. (1984) for a redefinition of these requirements as applied to $R*$).

Figure 3.5: A diagrammatic proof of Arrow's Theorem (Blackorby et al., 1984)

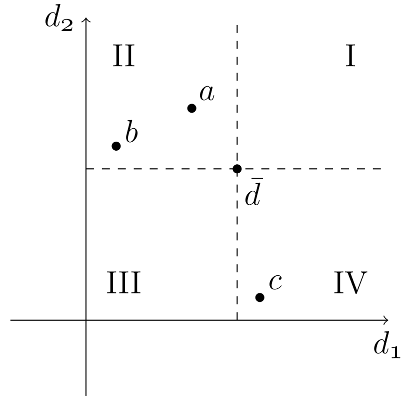

can now establish that all points in region II (or region IV) must be ranked the same way against $\bar{d}$. All points in region II are lower ranked by dimension 1 than $\bar{d}$ ($d_1 < \bar{d}_1$) and all points in region II are higher ranked by dimension 2 than $\bar{d}$ ($d_2 > \bar{d}_2$). Take a from Fig. 3.5 and assume $aP * \bar{d}$ (that is, a is strictly preferred to $\bar{d}$). There is an increasing monotone transformation[26] of $\bar{d}_1$ to itself and a_1 to b_1. This preserves dimension 1's ordinal ranking since we have made sure that a_1 is mapped to a smaller value than $\bar{d}_1$. Similarly, we can map $\bar{d}_2$ to itself and a_2 to b_2, preserving dimension 2's ranking.

Since the rankings of dimension 1 and dimension 2 are preserved, if $aP * \bar{d}$, then $bP * \bar{d}$. And because nothing hinges on the particular points a and b, we have established that all of region II is ranked in an identical fashion with respect to $\bar{d}$. Finally, we can eliminate the option that points in region II are ranked indifferently to $\bar{d}$ because, by the transitivity of indifference, $aI * \bar{d}$ and $bI * \bar{d}$ imply $aI * b$. But point a must be ranked above point b by the same reasoning we used to rank points in region I above points in region III (see Fig. 3.5).

Blackorby et al. (1984) then show that the ranking in region II must be opposite that of region IV. Suppose points in region II are ranked above $\bar{d}$. The idea is that there exists an increasing monotone transformation that will map a to $\bar{d}$ and $\bar{d}$ to c. Since such transformations preserve ordinal rankings, and since a is ranked above $\bar{d}$ by assumption, a remains ranked above $\bar{d}$ after the transformation—that is, $\bar{d}$ is ranked above c. Again, our choice of

26. In fact, there are infinitely many if all we care about is ordinal ranking, see (131).

Figure 3.6: Diagrammatic dictators

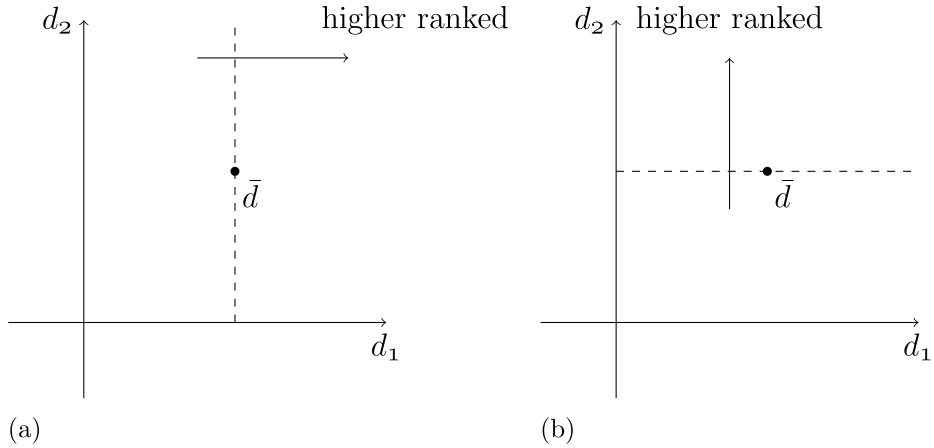

points was arbitrary, and we conclude that the points in region II must be ranked opposite to the points in region IV with respect to $\bar{d}$.

Finally, Blackorby et al. (1984) show that if two adjacent regions have the same ranking with respect to $\bar{d}$, the same ranking holds for any point on their common boundary. This leads to a dictatorship, as in Fig. 3.6 (Gaertner, 2009, 33). If the points in region II are ranked above $\bar{d}$, then the points in regions I and II and the points along their common boundary are ranked above $\bar{d}$. An alternative will be ranked higher in the aggregated ranking only if it is ranked higher according to the vertical dimension—dimension 2 in Fig. 3.6(b). Similarly, if the points in region IV are ranked above $\bar{d}$, then the points in regions I and IV and the points along their common boundary are ranked above $\bar{d}$. An alternative will be ranked higher in the aggregated ranking only if it is ranked higher according to the horizontal dimension—dimension 1 in Fig. 3.6(a).

Either dimension 1 or dimension 2 is a dictator.

As Gaertner (2009, 34) suggests, Blackorby et al.'s approach to Arrow's result "demonstrates the far-reaching consequences of [a] purely ordinal approach where [dimensional measures] are determined up to strictly monotone transformations." These consequences have special force in the semantics of multidimensional adjectives, where dimensional measures along one dimension (like blood pressure) are not comparable to dimensional measures along

Figure 3.7: Pairwise voting: one-pair and two-pair (Saari, 2008, 33).

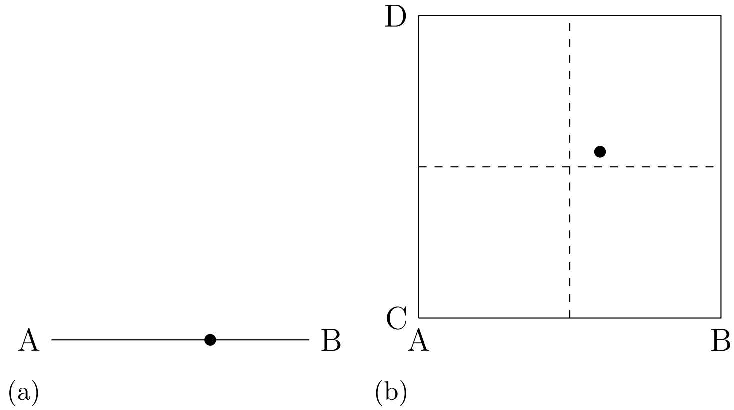

another (like heart rate) (Moltmann, 2005; Sassoon, 2013b). As Blackorby et al. (1984) shows, the aggregated measure of healthiness ($R*$ above) inherits and amplifies these informational constraints. Thus, desirable and fully justified limits on the information conveyed by adjectival dimensions point toward a paradox of choice.

Blackorby et al.'s formulation of Arrow's result reveals the importance of the multidimensional choice space in generating such paradoxes—the curse of dimensionality. The ranking of entities imposed by the multidimensional adjective *healthy* loses information associated with each dimensional ranking imposed by its component dimensions like blood pressure and heart rate.

To illustrate, we consider Saari's (2008) discrete voting example involving a choice between two alternatives, A and B, and a majority-vote aggregation function. If A receives 40 votes and B receives 60, the aggregation function will properly declare B the winner. Representing this result on a line between A and B, the point $\bullet$ in Fig. 3.7(a) represents both the election outcome and the profile of voters (how much support each alternative received from the voters).

Now suppose there are two pairs to decide between: $\{A, B\}$ and $\{C, D\}$, as in Fig. 3.7(b), with 55% supporting D over C (each corner in Fig. 3.7(b) represents a voter type, (A, C) or (A, D) or (B, C) or (B, D)). While the point $\bullet$ accurately describes the outcome of this election, it is compatible with at least two different profiles of voters. In one profile, 55%

of the voters prefer the B, D outcome, 40% prefer A, C, and 5% prefer B, C. In another, 45% prefer B, C, 40% prefer A, D, and only 15% prefer the actual outcome B, D. To put it another way, in the first profile, the majority supports the outcome, and in the second, 85% oppose it.

The domain of the aggregation function increases from a one-dimensional line in Fig. 3.7(a) to a higher-dimensional space in Fig. 3.7(b). In this higher-dimensional space, the outcome is compatible with a wider range of profiles (Saari, 2008, 35, fn. 7). While Saari frames this problem in terms of pairs of alternatives, his point holds for Blackorby et al.'s (1984) framing as well: in a multidimensional choice space, we lose the ability to reliably "unwind" result of the aggregation process to obtain meaningful information regarding the components of the domain of the aggregation process. In Fig. 3.7(b), it is no longer possible to determine how many votes each candidate received by looking at the outcome of the election; the aggregation process loses information as compared to Fig. 3.7(a). And as the number of alternatives increases, the information loss compounds. For 3 pairs, the electoral outcome is compatible with a mind-boggling five-dimensional domain of profiles.[27] In this case, we no longer have a guarantee than *any* individual voter preferred the actual outcome of the election.

3.7 Conclusion

On the degree-based account of the semantics of gradable adjectives, the degrees associated with gradable adjectives are formalized by a triple $\langle D, <, \delta \rangle$ including the set of degrees D, an ordering on this domain $<$, and a dimension δ that provides the property to be measured. For multidimensional adjectives like *healthy*, the relevant measured property (healthiness) is derived from multiple component dimensions in context, like blood pressure, cholesterol, and heart rate. Therefore, the degrees associated with multidimensional adjectives like *healthy* are associated with a triple $\langle D, <, f(\delta_1, \ldots, \delta_n) \rangle$, where f aggregates the scalar properties

27. As the dimensionality of the domain of profiles increases, the prevalence of potentially "degenerate" profiles also increases. In the voting example, this includes cyclic profiles (like $A > B$, $B > C$, and $C > A$).

associated with the component dimensions $(\delta_1, \ldots, \delta_n)$ into a single measurable property.

Arrow's theorem shows that seemingly sensible constraints on f—like Unanimity, Unrestricted Domain, Independence of Irrelevant Alternatives, and Non-Dictatorship—in a multidimensional choice space may lead f to fail in particular ways. Either the constraints themselves will be violated or other desirable properties like transitivity or continuity will be violated. Vagueness effects flow from these failures.

On this view, then, the intransitivity of the indifference relation results from the imposition of Arrowian constraints in a multidimensional choice space. In such spaces, the aggregation function—constrained by what appear to be reasonable assumptions—loses information. Independence in particular plays an information-losing role in the semantics of gradable adjectives. By hypothesis, each component dimension of a multidimensional adjective (like blood pressure for *healthy*) is a linear order: transitive and complete. If the aggregation function f of a multidimensional adjective obeys Independence, then the aggregation function is equivalent to one that considers only pairs of entities, like (127). But then the aggregation function f ignores the fact that the component dimensions are transitive. This lost information leads to an intransitive indifference relation associated with the multidimensional adjective, and an intransitive indifference relation is associated with vagueness effects.

This is the sense in which vagueness is semantic indecision: the standard approach to degree-theoretic semantics of multidimensional adjectives, plus reasonable assumptions about how dimensions are "aggregated" into one scale, yield intransitivities commonly associated with vagueness. The sorites premise is false but appears to be true because the aggregation procedure makes only pairwise judgments.

CHAPTER 4

MULTIDIMENSIONALITY FOR UNIDIMENSIONAL ADJECTIVES

Adjectives like *long*, *tall*, *expensive*—so-called "relative" adjectives—are vague (135), but they are not sensitive to contextually relevant dimensions (82b, 83b, 85b, repeated from above).

(135) Premise 1: A 3 hour wait for coffee is long.

Premise 2: Any wait 1 second shorter than a long wait is long.

Conclusion: A 5 second wait for coffee is long.

(82b) * The wedding is long {with respect to, in} temporal duration (but not with respect to space).

(83b) * The table is long in {all, most, three, some} respects.

(85b) * The table is long except with respect to temporal duration.

It's possible, of course, that vagueness effects in unidimensional adjectives and multidimensional adjectives stem from different sources, but this approach seems to miss a generalization.

Condorcet's paradox and the sorites paradox are two sides of the same coin. They both arise from a common source, the aggregation of many judgments into one. And in both cases, plausible constraints on the aggregation process lead to paradox in a multidimensional choice space. These constraints—generalized versions of Arrow's assumptions—are continuity (small changes in the input of the aggregation process should lead to small changes in the output), unanimity (where judgments agree, the aggregation process should reflect this agreement), and anonymity (no judgment is more important than another).

In a multidimensional choice space, continuity, unanimity, and anonymity are not mutually consistent. The sorites paradox and other vagueness effects result from this inconsistency.

Thus, the sorites paradox is paradoxical because it represents a choice constrained by inconsistent assumptions. Here, I extend the collective choice approach to adjectives like *tall*. I also discuss purportedly non-vague adjectives like *full*, and show how the "indecisional" approach accounts for the desiderata of a theory of vagueness. Finally, I conclude with a brief survey of this account's other advantages.

4.1 Many standards

Vague (unidimensional) adjectives are context-sensitive in more than one way. Norm-based adjectival standards (like the average height, or height of the majority, etc. for *tall*) are too rigid to support the boundarylessness of vague predicates (Kamp, 1975). For example, a person may be taller than average and yet not *tall* (Kennedy, 2007). Furthermore, a sorites paradox is possible for *tall* even if the heights of all the entities in the comparison class are known; for a norm-based standard alone, this is impossible.

Second, evidence from implicit comparisons suggests that adjectival standards incorporate more than one type of information. As Kennedy (2007) shows, implicit comparisons like (136a) are incompatible with "crisp judgments," judgments based on small but noticeable differences in degree. In (136a), the positive form *long* conveys the fact that there is an asymmetric ordering between two objects along some dimension, just like the explicit comparative form *longer* (136b). In particular, both (136a) and (136b) can be used to make a claim about a 100-page book in opposition to a 50-page book. However, implicit comparison is infelicitous in contexts requiring crisp judgments. For example, (136a), but not (136b), is infelicitous when used to make a claim about a 100-page book and a 99-page book.

(136) a. This book is long compared to that book.

b. This book is longer than that book.

As Kennedy (2007: 18) explains, "Since ... *long* is true of an object if and only if it stands out in length in the context of evaluation, an assertion of [(136a)] involves a commitment to the highly unlikely position that a difference of one page could actually be relevant to whether a book of the size of the two under consideration stands out in length or not." The semantics of *long* therefore takes into account more than a judgment about some cutoff point between *long* and *not long*; it also takes into account a "threshold" value between the *long* entities and the *not long* entities.[1]

This multicriterial approach to adjectival standards is mirrored in other analyses of vagueness phenomena. For instance, Kennedy's threshold semantics finds its reflection in the semi-orders of van Rooij's (2011) contextualist account and in Fara's (2002) "interest relative" account. Fara (2002: 64-65) writes that "the property expressed context-invariantly by *tall* is a property which is such that whether a thing has it depends not only on heights, but on other things as well," including "what our interests are." For yet another example, supervaluationist approaches rely on an interpretation function that simultaneously takes into account multiple adjectival cutoffs (Fine, 1975; Kamp, 1975; Sassoon, 2013b). According to these theories, an entity that exceeds all the relevant cutoffs for height counts as *tall,* an entity that exceeds none of the cutoffs counts as *not tall,* and an entity that exceeds some but not all cutoffs counts as neither *tall* nor *not tall.*

As approaches to vagueness phenomena, many-valued logics also confirm the multicriteriality of adjectival standards. For example, Cobreros et al. (2012) and Burnett (2014) develop a notion of "tolerant truth" in order to model vagueness effects like borderline cases, which (the authors argue) admit contradictions (137).

(137) Clarence is tall and Clarence is not tall.

1. Kennedy also points out that this effect is more general. For example, it appears when the positive form is used in a definite description.

(1) Pass me the long book.

(2) Pass me the longer book.

However, in these accounts borderline cases are only possible in models with three or more entities in the comparison class (Cobreros et al., 2012, 354). This suggests that vague predicates make use of distributional information from the comparison class. Indeed, experimental evidence suggests that classification under adjectives such as *tall* depends not only on entity's height but also on how the heights of other entities in the comparison class are distributed (Solt & Gotzner, 2012; Schmidt et al., 2009). Still others argue that metric distance from a prototype plays a role (Kamp & Partee, 1995; Douven et al., 2013).

For example, Lassiter & Goodman (2015) make the case that the shape of the probabilistic distribution of heights is an element of the semantics of *tall*. This approach draws on a picture of pragmatic feedback involving inferences alternating between a "literal listener", a speaker, and a "pragmatic listener." Lassiter & Goodman's model is "the interpretation process of a listener who uses literal interpretation as a base case and reasons to some finite depth. . . . The pragmatic listener L1 reasons about the utterance choices of a simulated speaker S1, who reasons about the interpretation of a literal listener L0, who does not reason pragmatically" (p. 7). The idea is that a listener updates her information state upon hearing some utterance by reasoning about alternative utterances and the probability that these utterances accurately describe some state of the world.

For example, Lassiter & Goodman describe a simulated interpretation of the utterance *Al is tall.* This utterance contains a free variable—the threshold value for what counts as *tall*—and the pragmatic listener's job is to infer the value of this free variable in a way that balances truthfulness and informativity. The literal listener L0 cares only about the truth of the utterance u relative to possible thresholds V, where A is Al's height (i.e., the answer to the question, "how tall is Al?").

(138) $\mathrm{PL}_0(A|u,V) = \mathrm{PL}_0(A|\ [\![\mathrm{u}]\!]^V = 1)$

The simulated speaker, S1, chooses a particular utterance u based on its informativity, which is in turn based on how the literal listener L0 would respond to it.

(139) $\mathrm{PS}_1(u|A,V) = \frac{\exp(\alpha \times \ln[\mathrm{PL}_0(A|u,V) - C(u)])}{\sum_{u' \in \mathrm{ALT}} \exp(\alpha \times \ln[\mathrm{PL}_0(A|u',V) - C(u')])}$

Here, (139) says that S1's choice of u conditioned on A and V is proportional to the utility (for the speaker) of u, calculated in terms its informativity to the literal listener L0.[2] The speaker's best choice is determined against what the speaker could possibly say—alternative utterances $u' \in \mathrm{ALT}$. Finally, the pragmatic listener L1 interprets u using Bayesian inference, assigning to each A and V a probability proportional to the product of (a) the probability that the speaker would have chosen to employ u if A were the true answer, or V the true threshold, and (b) the prior probability that A is true, or V the actual threshold.

(140) $\mathrm{PL}_1(A,V|u) \propto \mathrm{PS}_1(u|A,V) \times \mathrm{PL}_1(A) \times \mathrm{PL}_1(V)$

The result of this model, Lassiter & Goodman (2015, 20) reveal, is "a probabilistic 'sweet spot' interpretation for scalar adjectives which is highly sensitive to the statistical information encoded in the prior."[3] Tellingly, this model depends on the prior distribution of both heights and thresholds[4]—multiple types of information.

However the standard of *tall* is fixed, multiple types of information figure into the fixing. Instead of associating *tall*'s adjectival standard with one value based on height, I instead associate the standard with a vector of several values representing potentially different ways of ranking members of the comparison class. Each of these ranking methods may be considered a "standard" in a meaningful sense. Recapitulating the insight of supervaluationism, if multiple standards are available in context, then there are multiple ways of ranking members of the comparison class. And if this is true, then aggregating these multiple rankings will

2. The parameter $\alpha > 0$ determines how closely Lassiter & Goodman's stochastic choice approximates deterministic utility maximization. Informativity is defined in terms of "negative surprisal" to the literal listener minus the cost to the speaker, $C(u)$, of making the utterance.

3. The alternative utterances to *Al is tall* that Lassiter & Goodman consider are *Al is short* and the empty utterance $\emptyset$.

4. Lassiter & Goodman (2015, 19) assume a uniform distribution for thresholds. Their reasoning justifies the assumption of anonymity, discussed below.

give rise to paradoxes of choice in precisely the same ways that multiple rankings gave rise to paradoxes of choice in the semantics of multidimensional adjectives.

For example, the vectors in (77) represent different ways of ranking the members of the comparison class. In (141a), the "ordinal rank" standard would rank members of the comparison class based on the ordinal rank of their heights, while the "distance from prototype" standard would rank members of the comparison class based on their distance from some contextually determined prototype. Still another standard might rank the members of the comparison class relative to some relevant distribution of heights, as in Lassiter & Goodman (2015) (141b). This is what I mean by multidimensionality for unidimensional adjectives. I will refer to a particular vector of values $(a, b, \ldots)$ as a *profile.*

(141) a. (ordinal rank, distance from prototype)

b. (ordinal rank, distance from prototype, distribution, . . .)

The availability of multiple different standards in the same context accounts for the felicity of seeming contradictions like (137). On this view, (137) is really something like (142a).[5] Contrast this with (142b), which is a true contradiction. Notably, the example in (142a) uses the same "respect" locution as used with multidimensional adjectives.[6]

(142) a. In one respect, Clarence is tall, and in another respect, Clarence is not tall.

b. # In one respect, Clarence is tall, and in that same respect, Clarence is not tall.

If we have multiple standards, we need a way of aggregating them into a coherent choice. I reserve the function **s** for this role. As is clear from (143), **s** is an aggregation function.[7]

(143) $\mathbf{s}(\text{tall}) = f(\text{ordinal rank, distance from prototype, distribution, } \ldots)$

5. This may be an example of Ripley's (2011) "non-indexical contextualism" response to borderline case contradictions (see also Alxatib & Pelletier (2011)).

6. Barker (2002) would likely call uses like (142a) metalinguistic because they convey information about the prevailing standards of tallness but not much else (like Clarence's actual height, for example).

7. While the indecisional approach is committed to the existence of a function that aggregates rankings along multiple dimensions, this function need not be the standard function s itself. For example, it is possible

The semantic task for unidimensional adjectives is to aggregate the choices of multiple standards into a single standard. This becomes the "standard of comparison." Unidimensional adjectives are therefore collective choice functions, and they are subject to paradoxes of collective choice.

4.2 Topological collective choice

Generalizing Arrow's Theorem, Chichilnisky (1982) provides a topological proof of the paradox of collective choice. Saari (1997) and Baigent (2010) present informal versions of Chichilnisky's proof, discussed in chapter 3. Suppose Nino and Elena want to go camping along the shore of a perfectly circular lake. Nino and Elena may prefer the same geographic location along the shore, or they may not. If they agree on the location, that is where they will camp. Call this feature of the decision rule "unanimity." And the decision will not depend on who chose what: if Elena picks location 1 and Nino picks location 2, the outcome (whatever it is) will be the same as if Nino picked location 1 and Elena picked location 2. Call this feature of the decision rule "anonymity." Finally, the decision rule should be "relatively insensitive to small changes in individual preference" (Chichilnisky, 1982, 337). This last requirement brings a kind of stability into the decision rule (Lauwers, 2009). Call this feature "continuity."

Let S^1 denote the shore of the lake, let e be the location Elena picks, n the location Nino picks, and let $f(e, n)$ be the compromise choice, the result of aggregating Elena and Nino's preferences. Then a decision rule that chooses $f(e, n)$ such that $f(e, n)$ represents the unique shortest (arc) distance between e and n is both unanimous and anonymous. It is unanimous because, if $e = n$, then $f(e, n) = e = n$. It is anonymous because $f(e, n) = f(n, e)$.

However, it is not continuous. Holding e fixed, as in Fig. 4.1, if Nino's preferred location n moves continuously in a counterclockwise direction, so does $f(e, n)$—until, that is, n

to locate the aggregation function f within the triple associated with a unidimensional gradable adjective: $\langle D, >, f(\sigma_1, \ldots, \sigma i)\rangle$. Here, σ represents some ranking relevant to setting the standard, like ordinal rank.

Figure 4.1: Aggregation of location preferences along a circular lake.

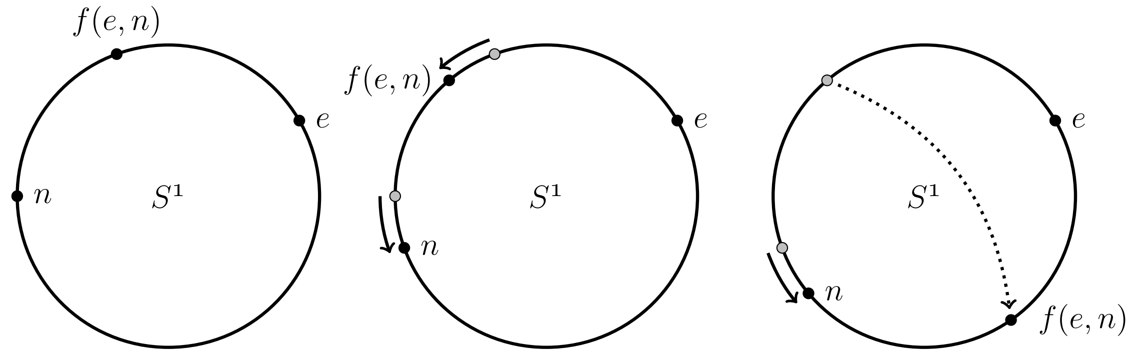

reaches the antipode of e, in which case $f(e, n)$ abruptly jumps to the other side of the lake. Chichilnisky's theorem generalizes this behavior.

Chichilnisky's theorem There is no continuous aggregation rule $f : S^1 \times S^1 \rightarrow S^1$ that satisfies unanimity and anonymity.

In terms of our example, there is no "fair" way to aggregate the location preferences along the shore of the lake. Any way we try to give Nino and Elena an equal say over the outcome, we will be stymied so long as we treat closely located camping spots similarly.

The camping problem and the sorites paradox are alike. To construct an analogy between Chichilnisky's social choice paradox and the sorites paradox, it remains to show that the space of profiles is analogous to a circular lake, and that adjectives obey the relevant analogues to unanimity, anonymity, and continuity.

First, it is possible to interpret adjectival standards as representing preferences in a circular (or spherical) choice space. An aggregation of these standards constitutes a collective choice subject to Chichilnisky's result.

Take the choice space X to be a subset of $\mathbb{R} \times \mathbb{R} \times \ldots \times \mathbb{R} = \mathbb{R}^n$. In Chichilnisky's native habitat of consumer theory, the elements of the choice space may be bundles of goods, like 4 bottles of wine and 3 bottles of beer, represented by the ordered pair $(4, 3)$. In our interpretation, the elements will be relevant thresholds along multiple potential standards. Recall (a modified version of) the possible vector for the threshold of *tall* (141b).

(141b) $\mathbf{s}(\text{tall}) = \langle\text{height norm, ordinal rank in comparison class}\rangle$

An ordered tuple $\langle 4, 3\rangle$ in this context would stand for a (distance from the) height norm of 4 (in relevant units) and an ordinal rank of 3 in the comparison class. These numbers represent a potential threshold along each of the relevant standards between the extension and anti-extension of a gradable adjective, and the vector $\langle 4, 3\rangle$ represents a point in a multidimensional choice space composed of multiple standards.

To interpret preferences over X—to answer what it means to prefer one cutoff over another—we may consider any reasonably consistent (i.e., transitive, complete, and continuous) choice rule.[8] This choice rule (a ranking of some sort imposed by the collective choice function f) corresponds to an adjectival standard. For example, following Lassiter & Goodman (2015), a probability-based standard may assign different probabilities to different cutoffs $\vec{v_1}$ and $\vec{v_2}$. Similarly, different norm-based standards (like a mean-based standard and a mode-based standard) may assign different values to $\vec{v_1}$ and $\vec{v_2}$ depending on their (absolute or squared) distance from the norm, just as a prototype-based standard may assign different values to $\vec{v_1}$ and $\vec{v_2}$ depending on their closeness in conceptual space to some prototype (Douven et al., 2013). And there may be standards that represent some combination of judgments—for instance, a standard that combines norm-based and threshold-based determinations.[9] The following table illustrates the analogy between aggregating adjectival profiles and Chichilnisky's setup.

(144)

	camping	Chichilnisky	adjectival semantics
elems. of choice space	locations	alternatives	profiles
choosers	campers	voters	standards

8. Chichilnisky used a utility function $u : X \to [0, \infty)$ that assigns to elements of X some "utility." A point $x \in X$ is preferable to a point $y \in X$, written xRy, iff $u(x) \geq u(y)$.

9. Such a standard might take into account threshold values only for those cutoffs that reach the norm, for example.

Different standards will prefer different profiles.[10] For example, a norm-based standard may rank the $(5, 4)$ profile equivalently to the $(5, 3)$ profile for the simple reason that the norm-based standard only cares about mean height, not relative position in the comparison class. A different standard—one that takes into account position in the comparison class—would not rank these profiles equivalently. These equivalences generate (linear) indifference sets or equivalence classes, and the "gradient vector" of these indifference sets indicates the direction of greatest increase along the relevant standard, as $\vec{v_{s_1}}$ and $\vec{v_{s_2}}$ in Fig. 4.2(a) and (b). In other words, the gradient vector points in the direction of profiles ranked highest by the relevant standard.

We are interested in choosing one point in this multidimensional space from among the available options. Following Chichilnisky, we consider ordinal preferences among the standards as a simplifying assumption. (The introduction of cardinal preferences does not avoid Chichilnisky's result, but I assume that we are interested in a choice of standard, not a cardinal ranking of standards.) Therefore, we normalize these gradient vectors to length 1, ignoring preference strength. Now, without loss of generality, we can lift these vectors and place them at the origin of a unit circle S^1 (Fig. 4.2(c)). This makes it easy to see that each point on S^1 represents a standard—namely, the standard whose unit vector points to that location on S^1 (Lauwers, 2009).

An aggregation rule is a map F (as in 145) from a tuple of standards to a single standard.

(145) $F : S^1 \times S^1 \times \ldots \times S^1 \longrightarrow S^1$

This recalls the camping problem (Fig. 106). The set of standards involved in the interpretation of an adjective like *tall* can be represented as points along the boundary of a circle, just like camping spots along a lakeshore.

Accordingly, if the map F is continuous, unanimous, and anonymous, then F does not exist. Formally, the three assumptions unanimity, anonymity, and continuity look like (146).

10. Indeed, Chichilnisky's result only works for suitably diverse voters. Adjectivally speaking, the failure of diversity results in absolute adjectives. See section 4.2.2.

Figure 4.2: Gradient vectors from (a) and (b) in a unit circle.

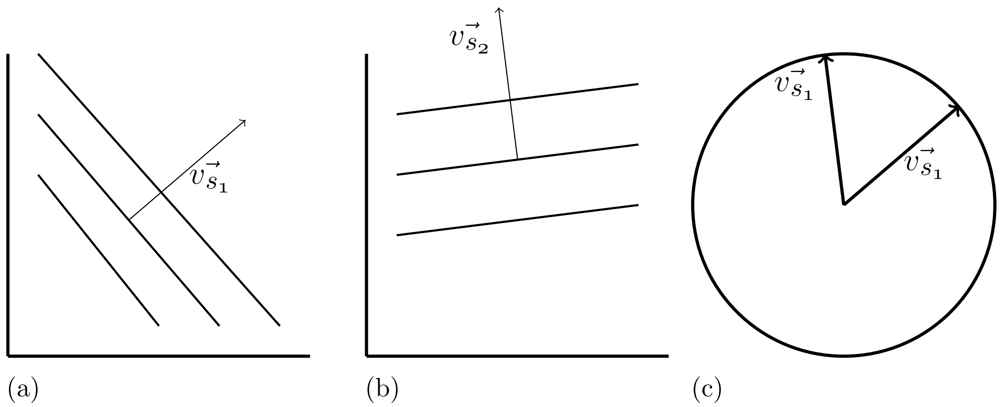

(146) a. **Unanimity:** for each point $p \in S^1$, $f(p, p, \ldots, p) = p$.

b. **Anonymity:** for any $p, q \in S^1$, $f(p, q) = f(q, p)$.

c. **Continuity:** for any map $f : \mathbb{R}^m \to \mathbb{R}^k$, the map f is continuous at $p \in \mathbb{R}^m$ if

i. for each $\epsilon > 0$, there exists a $\delta > 0$ such that $|f(q) - f(p)| < \epsilon$ as soon as $|q - p| < \delta$.[11]

Unanimity and anonymity appear to be reasonable constraints on the choice of adjectival standard. For example, unanimity implies that if every relevant standard would rank x as tall, then x is tall.[12] Similarly, if every relevant standard would judge x as not tall, then x is not tall. Anonymity reflects the idea that there is no contextually relevant standard that is somehow more important than another. For example, Lassiter & Goodman (2015, 19) justify anonymity for standards: "The assumption of uniform priors on semantic variables

11. This is only one possible definition of continuity consistent with the theorem. The definition of continuity is dependent on the particular topology imposed on the space of standards (that is, what constitutes the "open" sets).

(1) **Continuous map**: Let X and Y be topological spaces. A map $f : X \to Y$ is called continuous if the inverse image of opens sets is always open. (Janich, 1984)

The particular topology need only reduce to the Euclidean topology, and "any reasonable topology, certainly any topology which has been used on preference spaces, satisfies this condition" Heal (1997, 3).

12. This is analagous to "supertruth" in supervaluationism (Williamson, 1994).

means that . . . all possible thresholds for *tall* are equally good candidates a priori. . . . This assumption seems to be justified at least in the case of scalar adjectives, where non-uniform priors would limit the flexibility of interpretation: in other words, [standards for *tall*] were strongly biased toward human-like heights, this bias would influence the interpretation of *tall skyscraper* in strange ways. We eliminate this possibility by employing uniform priors." In the context of collective choice, the aggregation rule is invariant under permutations of the incoming standards. At least for relative adjectives, this seems like an uncontroversial assumption.

As discussed in chapter 3, continuity is implied by tolerance. In topological social choice, continuity is sometimes explained as "stability." For example, if Nino suddenly changes his opinion and claims that his favorite camping site is a site next to his previous choice, then the output of the aggregation rule should change at most to a site next to the previous one. As Lauwers (2009, 454) explains, "Continuity is a form of stability of the social outcome with respect to small changes in the individual preferences."

4.2.1 Vagueness with holes

And just like the lakeshore, the space of adjectival standards has a "hole" in it. In the camping problem, the campers' choices were limited to the shore of the lake—they could not decide to camp *in* the lake. This lacuna in the domain of the aggregation function is topologically a hole. Similarly, in the present example, the space of adjectival standards consists of the boundary of the unit circle S^1. It is the hole-y structure of the space of standards that is ultimately responsible for Chichilnisky's result and, consequently, vagueness effects. If the domain of the aggregation function contains holes, then unanimity, anonymity, and continuity are mutually incompatible. Vagueness effects result from this incompatibility.

In order to evaluate this incompatibility, we need the tools of algebraic topology, which is handy for studying continuous maps between different kinds of spaces. Topology doesn't care about metrics or distances between points in a space, but rather the overall "shape" of a space. For example, a topologist does not distinguish a circle from a square because it is possible to construct a continuous bijective map (whose inverse is also continuous) from one to the other. This is a fancy way of saying that one can be stretched into the shape of the other.

But there is no way to stretch a space with a hole in it, like a circle, into a space without a hole, like a disk. Topological spaces are equivalent if it is possible to continuously deform one into another. This notion of continuous deformation goes by the name "homotopy." And there is no way to continuously deform a hole-y space into a space without holes. Below, I rely heavily on Lauwers (2009), but other treatments can be found in Chichilnisky (1980, 1982); Baigent (2010); Lauwers (2000); Mehta (1997).

Let $\mathcal{P}$ be a topological space of individual preferences, like S^1. We are looking for an aggregation rule $F : \mathcal{P}^n \to \mathcal{P}$ that satisfies continuity, unanimity, and anonymity. The unanimity restriction on F means that $F(p, p, \ldots, p) = p$, i.e. if every individual has the same preferences, the aggregation function returns these preferences. Call the set of unanimity profiles DIAG, and define Δ to be the unanimity restriction on F, $\Delta(p, p, \ldots, p) = p$. The question is then whether there is a unanimous, anonymous, and continuous $F : \mathcal{P}^n \to \mathcal{P}$ such that $F \circ i = \Delta$, where i is the inclusion map of the unanimity profiles into the set of profiles.

(147)
$$\begin{tikzcd} \mathcal{P} \times \mathcal{P} \times \ldots \times \mathcal{P} \arrow[dr, "F"] & \\ \mathrm{DIAG} \arrow[u, "i"] \arrow[r, "\Delta"'] & \mathcal{P} \end{tikzcd}$$
(Lauwers, 2009, 460)

The answer is no, and to understand why the requirements of unanimity, anonymity, and continuity conflict, we translate this topological question into an algebraic one. The translation takes topological spaces into group theoretic objects in a principled way. In particular, spaces are translated into groups, and a continuous map between spaces is translated into a homomorphism between groups. (This is done through a functor called the first homotopy group.)

Briefly, a group $(G, *)$ is a mathematical object consisting of a set G and an operation $*$ that meets certain conditions (148).

(148) a. Identity: There is a special element e with $e * x = x * e = e$.

b. Inverse: For any x, there is an x^{-1} with $x * x^{-1} = x^{-1} * x = e$.

c. Associativity: For any x, y, z, $x * (y * z) = (x * y) * z$

For example, the set of integers and the operation of addition $(\mathbb{Z}, +)$ is a group under this definition. A homomorphism $\phi : G \to H$ between groups $(G, *)$ and $(H, \cdot)$ is a map that preserves algebraic structure, i.e. it takes the identity element in G to the identity element in H, the inverses in G to the inverses in H, and it respects the operation of $*$ as in (149).

(149) $\phi(g_1 * g_2) = \phi(g_1) \cdot \phi(g_2)$

We translate the space of individual preferences $\mathcal{P}$ into a group Π and the continuous function F into a homomorphism F_*, as in (150).

(150) $F_* : \mathit{\Pi} \times \mathit{\Pi} \times \ldots \times \mathit{\Pi} \longrightarrow \mathit{\Pi}$

This is a homomorphism from the group consisting of n-tuples of individual preferences $(\Pi \times \Pi \times \ldots \times \Pi)$ to a single preference (Π). The group analogue of the cartesian product is the "direct product," denoted $\times$, which takes two groups (say G and H) and constructs a new group $(G \times H)$. The elements of the new group are (g, h) with $g \in G$ and $h \in H$, and the binary operation $\times$ is defined component-wise, i.e. $(g_1, h_1) \times (g_2, h_2) = (g_1 * g_2, h_1 \circ h_2)$. Therefore, we note that for each a in Π, we have

(151) $(a, a, \ldots, a) = (a, e, \ldots, e) * (e, a, e, \ldots, e) * \ldots * (e, e, \ldots, e, a)$

with e the identity element.

Applying the homomorphism F_* to (151), unanimity and anonymity imply

(152)
$$\begin{aligned}
a &= F_*(a, a, \ldots, a) && \text{unanimity} \\
&= F_*((a, e, \ldots, e) * (e, a, e, \ldots, e) * \ldots * (e, e, \ldots, e, a)) && (151) \\
&= F_*(a, e, \ldots, e) * F_*(e, a, e, \ldots, e) * \ldots * F_*(e, e, \ldots, e, a) && \text{homomorphism}
\end{aligned}$$

Figure 4.3: Loops on a closed disk and an annulus.

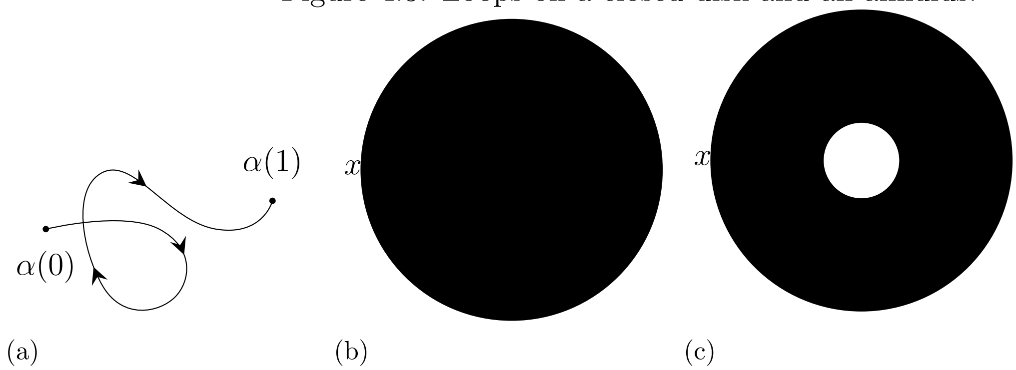

$$= F_*((a, e, \ldots, e) * F_*(a, e, \ldots, e) * \ldots * F_*(a, e, \ldots, e)) \qquad \text{anonymity}$$

Let $F_*(a, e, \ldots, e) = a'$. Since the group $(\Pi, *)$ is commutative (which I won't show), it turns out that we can treat the operator $*$ like we would treat addition $(+)$ over the integers. This means that

$$\begin{aligned} (153) \quad a &= F_*(a, a, \ldots, a) = a' + a' + \ldots + a' \\ &= na' \end{aligned}$$

And, in general,

$$(154) \quad F_*(a_1, a_2, \ldots, a_n) = n^{-1}(a_1 + a_2 + \ldots + a_n)$$

What (154) tells us is that the unanimous, anonymous, and continuous aggregation function we are looking for is something like arithmetic mean. But as the camping example illustrated, the mean is not continuous in certain spaces—namely, spaces with holes.

To prove this, we need the notion of the first homotopy group. A "path" is a continuous function $\alpha : [0, 1] \to X$ from the unit interval to a topological space X, with $\alpha(0)$ mapping to the initial point of the path and $\alpha(1)$ mapping to the final point (Fig. 4.3(a)). A "loop" is a path in which $\alpha(0) = \alpha(1)$. Usually, this point is identified at x_0, as in Fig. 4.3(b).

Roughly speaking, two paths are considered to be "homotopic" if they can be continuously deformed—shrunk, stretched, twisted, but not torn—into one another. The homotopy relation is an equivalence relation. It is reflexive (each path is homotopic to itself), symmetric (if α is homotopic to β, then β is homotopic to α) and transitive (if α is homotopic to β, and β is homotopic to γ, then α is homotopic to γ). We can use this relation to create equivalence classes of loops. Denote the equivalence class of the loop α as $[\alpha]$. Then, the "first homotopy group" is a group formed from the set of equivalence classes of loops starting at some point x_0.

In Fig. 4.3(b), for instance, the two loops α and β are homotopic to one another—we can continuously shrink the closed disk until α and β overlap completely. Moreover, we can keep shrinking the disk until both α and β are equivalent to the "constant loop," the loop that begins and ends at x_0 and goes nowhere in between (that is, the constant loop is the point x_0). And we can do this for all loops in the closed disk. Therefore, the set of equivalence classes of loops starting at x_0 in the closed disk is just the set of equivalence classes of the constant loop.

How do we make a group out of the set of equivalence classes of loops? We do this by defining an operation that respects the axioms of groups (148). This turns out to be the operation of path composition. The path composition operation ($\circ$) is constructed as follows: $\alpha \circ \beta$ is a loop that first goes around α and then around β.[13]

Imagine that the outer radius of the annulus is fixed, and that the inner radius increases until it equals the outer radius. Then we have a circle S^1 with the same homotopic properties as the annulus. The space of adjectival standards S^1 (or the sphere S^n) is homotopic to the annulus is Fig. 4.3(c).

Now we are in a position to define the first homotopy group. This is the group $(\pi_1(X, x_0), \circ)$, where $\pi_1(X, x_0)$ represents the set of equivalence classes of loops in a space X starting at x_0, and $\circ$ is the path composition operation. In the group $(\pi_1(D, x_0), \circ)$, where D represents the closed disk, each closed path is homotopic to the constant loop. Therefore, this group is

13. The identity element is the constant loop, denoted $[\alpha_0]$, and the inverse element for any loop α is the same loop traveling in the opposite direction.

trivial, i.e. it is isomorphic to the group consisting only of zero and the addition operation $(\{0\}, +)$.

However, consider the annulus in Fig. 4.3(c). Loops in the annulus cannot be shrunken to a point because of the hole in the middle. In this case, loops are homotopic just in case they go around the hole in the same direction the same number of times (like α and β in Fig. 4.3(c)). One can walk along the path α one time, or two times, or three .... And one can walk along the path α in the opposite direction once, twice, etc. In the end, the homotopy group of the annulus is isomorphic to the integers under addition, $(\pi_1(A, x_0), \circ) \cong (\mathbb{Z}, +)$, where A represents the annulus.

Imagine that the outer radius of the annulus is fixed, and that the inner radius increases until it equals the outer radius. Then we have a circle S^1 with the same homotopy group as the annulus, $(\pi_1(S^1, x_0), \circ) \cong (\mathbb{Z}, +)$.

We can now prove Chichilnisky's theorem and its linguistic corollary. If we interpret the circle S^1 as representing the set of standards a speaker may use to describe whether something is, for example, *expensive*, then a slight change in the aggregated standard represents movement along the boundary of the circle. If we imagine the standard changing (in a continuous manner) all the way around the circle, this is equivalent in group-theoretic terms to the integer 1. Suppose there is a unanimous, anonymous, and continuous function $F_* : \mathbb{Z}^n \to \mathbb{Z}$ (recall that the homotopy group of the circle S^1 is equivalent to the group $(\mathbb{Z}, +)$). By (153), repeated below, $1 = na'$.

(153) $1 = F_*(a, a, \dots, a) = a' + a' + \dots + a' = na'$

For $a = 1$, this means that there must be some integer a' such that $na' = 1$. Since $n \geq 2$ (representing the assumption of multiple potential standards), such an integer a' does not exist.

Thus, there is no aggregation rule F from a tuple of standards to a single threshold. Put still another way, "[W]henever there are holes in a domain, then some variable or agent dominates" (Saari, 2008, 28). Returning to the camping example, if two agents have antipodal location preferences, then there is no decision unless one agent gets at least a little more of what she wants. In the adjectival context, this results in discontinuities and hence vagueness effects, as explained in section 3.5.

4.2.2 Absolute adjectives

Further evidence for the social choice approach comes from non-vague adjectives, so-called "absolute" adjectives like *full*. On the present account, the property distinguishing absolute from relative adjectives is the dimensionality of their standards. If vagueness is a problem of collective choice, this distinction explains both why vagueness effects do not generally arise in absolute adjectives and why such effects may arise if more dimensions are introduced into the aggregation process.

Absolute adjectives have "rigid"—context independent—standards, while relative adjectives tend to have standards that vary with context. In degree-theoretic terms, Kennedy & McNally (2005) analyze absolute adjectives as describing entities that either possess a non-zero amount of the property indicated by the adjective, or they possess the maximal amount. These interpretations affect the types of default entailments the adjectives have. In (155a), for example, the door possesses zero degrees of "openness" (due to the interaction with negation) and therefore the continuation with *but it is still ajar* is semantically odd. In (155b), in contrast, the possession of zero degrees of "tallness" (not to say zero degrees of height) does not conflict with above-average height.

(155) a. # The door is not open, but it is still ajar.

b. Sam is not tall, but he is above average in height.

Kennedy (2007) provides four ways of distinguishing absolute adjectives from adjectives like *tall*: natural precisifications, entailments (like 155a), interpretations with the definite article (like *the straight one*), and participation in the sorites paradox. The latter is most relevant here. Absolute adjectives do not easily give rise to the sorites paradox because the second premise is judged to be false (156).

(156) Premise 1: A theater in which every seat is occupied is full.

Premise 2: Any theater with one fewer occupied seat than a full theater is full.

Conclusion: Therefore, any theater in which half of (none of, etc.) the seats are occupied is full. (Kennedy, 2007, 30)

This (along with Kennedy's three other types of evidence) supports the claim that absolute adjectives are distinct from relative adjectives. In particular, (156) suggests that absolute adjectives are not vague.

Despite the evidence in (156), however, absolute adjectives may sometimes exhibit vagueness effects. In some cases, what looks like vagueness actually results from "imprecision." This is the type of effect responsible for an utterance like (157) made in the context of a theatre with very few people in it. This effect is often referred to as "loose talk."

(157) This theatre is empty.

In other cases—more troubling for the absolute vs. relative generalization—absolute adjectives are vague for reasons not easily traceable to imprecision. In (158), for instance, it is possible that the glass is filled only half-way (Rotstein & Winter, 2004).

(158) This glass of wine is full.

Examples like (158) aren't simply instances of loose talk; "it is difficult to imagine that we are taking the glass's maximal capacity as the standard" and speaking loosely (McNally, 2011, 158).

Since even absolute adjectives may be vague, various researchers have attributed the difference between absolute and relative adjectives to something other than the adjectival standard (the semantic locus of vagueness phenomena). This includes distinctions premised on rule- vs. similarity-based classification (McNally, 2011), counterpart vs. extensional comparison classes (Toledo & Sassoon, 2011), maximal vs. restricted comparison classes (van Rooij, 2011), and asymmetric vs. symmetric indifference relations (Burnett, 2014). These accounts share a common intuition, which I likewise adopt: in context-dependent cases like (158), the standards of absolute adjectives incorporate more information than they normally do. For (158), this might include information about the relevant size and shape of the glass or the purpose of the wine (e.g., for drinking or tasting). In McNally's (2011) version, for example, while the context-independent use of *full* only makes reference to volume, the context-dependent use in (158) requires taking into account both volume and the "type of container."

This approach is natural if the context-independent use (the use in 156) is the default use, as nearly all of the above accounts assume. For example, McNally's description realizes the difference between the context-independent and the context-dependent use of *full* in a manner equivalent to the distinction between (159a) and (159b), respectively.

(159) a. $\mathbf{s}$(full) = (volume)

b. $\mathbf{s}$(full) = (volume, type of container)

As the distinction between (159a) and (159b) implies, the vagueness of *full* is a function of the dimensionality of its standard. In their non-vague uses, absolute adjectives have the structure of a one-dimensional space of standards. These adjectives are seemingly sensitive to only one value, like volume in the case of *full* or degree of bend in the case of *straight*.[14] And if the space of standards is one-dimensional, the domain of the aggregation function is more like a line than a circle. Aggregation problems do not arise; it is possible for an aggregation function to be continuous, anonymous, and unanimous (Chichilnisky, 1982, 347).[15]

14. As a corollary, these adjectives are insensitive to variations in other dimensions. For example, van Rooij (2011) describes how absolute adjectives are insensitive to changes in the relevant comparison class.

15. This idea is contained within a theorem by Chichilnisky & Heal (1983).

Chichilnisky and Heal 1983. The space $\mathcal{P}$ allows for topological aggregation if and only if each closed path is homotopic to the constant path. (In other words, there are no "holes" in $\mathcal{P}$.)

Thus, absolute and relative adjectives are distinguished by the dimensionality of their standards. In multidimensional choice spaces, like the space of relative adjectival standards, unanimous, anonymous, and continuous aggregation functions fail. In one-dimensional choice spaces, in contrast, such aggregation rules are well-behaved.[16]

If the dimensionality of the standard distinguishes absolute and relative adjectives, then this is also a natural place to locate the source of vagueness effects. In the default case, absolute adjectives are distinguished from relative adjectives by measuring only one value, like volume in the case of *full*. In their vague uses, however, the standards of absolute adjectives (159b) look a lot like the standards of relative adjectives (141a).

(159b) $\mathbf{s}$(full) = (volume, type of container)

(141a) $\mathbf{s}$(tall) = (height, ordinal rank in comparison class)

An aggregation function defined over the multidimensional choice space of (159b) creates collective choice problems deriving from the same source as with relative adjectives: the mutual incompatibility of unanimity, anonymity, and continuity.

16. Perhaps some evidence for dimensionality as the distinguishing factor comes from multidimensional adjectives. Some multidimensional adjectives pattern like absolute adjectives with respect to modification (1) and definite descriptions (see Grinsell (2013)).

(1) a. Stephen is {perfectly, almost, ??slightly} healthy.

b. Tony is {??perfectly, slightly} abnormal.

c. Ruth is {perfectly, slightly} familiar with these routines.

But adjectives like *healthy*, *abnormal*, and *familiar* display vagueness effects. In particular, multidimensional adjectives like *healthy* have borderline cases and they participate in the sorites paradox. What distinguishes an adjective like *full* (in its default use) from *healthy* is not the presence or absence of scalar endpoints, but rather how much information the standard takes into account.

Figure 4.4: A Voronoi diagram (Douven et al., 2013, 142) and a collated Voronoi diagram (Douven et al., 2013, 146).

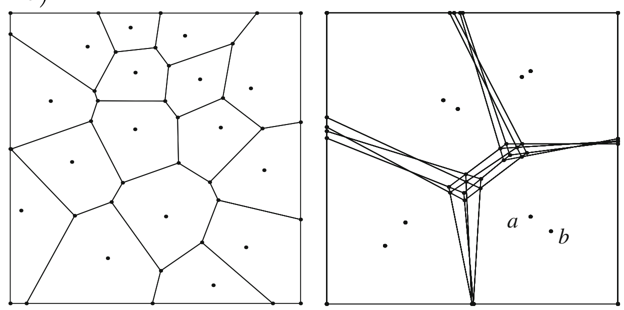

(a) (b)

4.3 Conclusion

One advantage of the topological approach is that it connects theories of linguistic meaning to psychological theories of concept representation. For example, Douven et al. (2013) present a theory of vagueness in terms of "conceptual spaces," geometric representations of concepts. Drawing on Gärdenfors (2004), these authors envision concepts as regions of conceptual space (unidimensional or multidimensional) over which a distance metric is defined. In order to represent concepts like *red*, Gärdenfors then includes a "prototype," which Douven et al. (2013, 141) describe as "a typical case, where this is understood as a case that competent members of a given language community in normal circumstances classify, or would classify, unhesitatingly as being typical." This combination results in a division of conceptual space into cells centered at the relevant prototype. The distance metric determines which points lie in which cell by determining which points are closest to each prototype. The product of this process is a "Voronoi diagram." In Fig. 4.4(a), the center of each cell is a prototype, and each cell consists of those points which are closer to its own prototype than to any other prototype.

Gärdenfors acknowledges it is unlikely that a concept like *red* has a unique prototype; rather, many shades of red may be typical instances of red (Douven et al., 2013, 143). In order to accommodate this insight, Gärdenfors replaces the assumption of prototypical "points" with prototypical "areas." In particular, Gärdenfors assumes that prototypical areas are circles, and that the distance between any two points on the circle is 0.

Gärdenfors's representation is very close to the one presented here. He represents concepts as topological spaces (metric spaces are topological spaces) with an analogue of holes. In his theory, the distance from a point in the "hole," the prototypical area, to the center of the space is imaginary.[17] As Douven et al. (2013, 144) note, this is "not easy to interpret if distances are supposed to reflect degrees of similarity." In the present account, there is no need for a distance metric, thereby avoiding this difficulty.

Douven et al. recognize this drawback and adopt an alternate strategy: each prototypical point determines a separate Voronoi diagram, and these diagrams collectively determine the conceptual space (Fig. 4.4(b)). They then develop a notion of a borderline case that is associated with points that fall along "thick" borders in a collated Voronoi diagram. This sort of approach makes great strides in uniting linguistic theories of vagueness (in particular data about borderline cases) with psychological representations of concepts. And one advantage of this unification, as Douven et al. make clear, is that it renders hypotheses about vagueness phenomena susceptible to familiar experimental testing.

Douven et al. acknowledge that their theory of vagueness must yet account for more phenomena.[18] In particular, it works well for perceptual cases but not complex or abstract cases (p. 138), and it has yet to be extended to "soriticality" (*id.*), concepts like *tall, poor* and *beautiful* (p. 141), and the boundarylessness (or higher-order vagueness) of both prototypical and borderline cases (p. 151-52). The present account, though not exactly in the same vein, makes progress on these extensions.

17. It is therefore doubtful that these spaces are convex in any meaningful sense, as Gärdenfors intends.

18. Both Douven et al. (2013) and the present approach rely on multiple standards. Douven et al. use multiple prototypes in their construction of the conceptual space for a vague adjective like *red.* These prototypes (or, more specifically, some distance from them) might reasonably be associated with the adjectival standard for *red.* Douven et al.'s insight is that vagueness effects become easier to capture if it is assumed that multiple standards factor into the meaning of a vague predicate simultaneously. The present account respects this insight.

A second advantage of the topological approach is that it provides an explanation for the promiscuity of vagueness effects across syntactic and semantic categories. As van Rooij (2011) observes, vagueness effects appear in common nouns (*heap*), adverbs (*very*), quantifiers (*many*), verbs (*start*), proper names (*Chicago*), and definite descriptions (*the border between Illinois and Iowa*). These parts of speech all involve multidimensional decisions—whether something is a heap, for instance, depends on its width and height, among other things—and therefore these parts of speech are all susceptible to paradoxes of decision making resulting in vagueness effects.

For example, the event-in-progress reading of the English incremental-theme progressive, as in (160), displays vagueness effects like borderline cases and the sorites paradox.

(160) a. Diana Nyad was swimming from Cuba to Florida.

b. Mary was wiping out the Roman army. (Landman, 1992, 18)

As Condoravdi (2009, 14)14 has observed, "[T]he truth of [*Mary was wiping out the Roman army*] depends on ... how much of the Roman army Mary has destroyed at the reference time in the world of evaluation." That is, the slaughter of 10,000 soldiers probably licenses the progressive, but the slaughter of dozens probably does not. What about the slaughter of 300? Or 500? The lingering uncertainty in these latter scenarios is the essence of a borderline case. Similarly, a sorites paradox for the progressive is also readily constructed, as in (161):

(161) **Premise 1.** (At the reference time in the world of evaluation, Mary had killed 10,000 soldiers) Mary was wiping out the Roman army.

Premise 2. Any event of "wiping out the Roman army" in which one less soldier was killed is still an event of wiping out the Roman army.

Conclusion. (At the reference time in the world of evaluation, Mary had killed 1 soldier) Mary was wiping out the Roman army.

An analysis that relativizes the interpretation of the progressive to multiple views of what constitutes the "normal course of events" accounts for these effects (Grinsell, 2014). In a modal theory of the progressive (Dowty, 1979; Landman, 1992; Portner, 1998; Condoravdi, 2009), the progressive is true if, in the normal course of events, the event would have reached completion. Grinsell (2014) argues that what constitutes the "normal course of events" is in fact determined by multiple different ways of ranking possible worlds. These multiple rankings of worlds are aggregated into one, triggering vagueness effects in the same way that multiple dimensional rankings trigger vagueness effects in multidimensional adjectives.

Phenomena like the vague progressive pose particular challenges to accounts of vagueness like those of Qing & Franke (2014); Lassiter & Goodman (2015). Approaches like Lassiter & Goodman's (2015) rely on the probabilistic distribution of the relevant property (e.g., height in the case of *tall*) and the possible values for the "threshold" variable θ. Lassiter & Goodman (2015), for example, make the case that the probabilistic distribution of heights is an element of the semantics of *tall.* This approach draws on a picture of pragmatic feedback involving inferences alternating between a "literal listener", a speaker, and a "pragmatic listener". The listener updates her information state upon hearing some utterance by reasoning about alternative utterances and the probability that these utterances accurately describe some state of the world, a process that depends on finding "a probabilistic 'sweet spot' interpretation for scalar adjectives which is highly sensitive to the statistical information encoded in the prior." Or as Qing & Franke (2014, 39) puts it, the possible values of the adjectival threshold relies on a probabilistic distribution of such values arising from "conventional semantic knowledge within a linguistic community."

Yet this conventional knowledge may sometimes be absent. With respect to progressives like (160b), it is unclear, for instance, whether conventional semantic knowledge includes a distribution of "wiping out" events. Indeed, it is unclear what such a distribution would distribute over: the duration of the event, how much damage was inflicted during the event, the proportion of "completeness" of the event (somehow defined), all of these or none? Speakers know these terms are vague likely without knowing the related property's probability distribution.

However, as intimated throughout this chapter and chapter 3, the analysis of vagueness effects in terms of limitations of collective decision making—including the topological approach—is heavily dependent on assumptions about how *diverse* the space of alternatives is. The approach to absolute adjectives presented in subsection 4.2.2 demonstrates this point. If the space of adjectival standards is one-dimensional rather than multidimensional, the paradoxes of choice and hence of vagueness will not arise.

The problem is in fact even more acute. Paradoxes of choice arise with some profiles but not with others. As Sen (1970) shows, collective choice does not always entail intransitivity. If the relevant dimensions agree on rankings of members of the comparison class, or if the relevant adjectival standards agree on profiles, then it's possible that collective choice yields a transitive aggregated ranking. Equivalently, if voters tend to agree on the relevant alternatives, the collective decision will not cycle. Only if the dimensions differ, or the standards differ, or the voters differ sufficiently will intransitivity or discontinuity result.

Sen's formal distinction is between so-called "multi-peaked" rankings and "single-peaked" rankings. Intuitively, single-peakedness limits the breadth of disagreement that various rankings can have over the relevant alternatives. For example, in Figure 4.5(c), each line displays only one peak. Interpreting Figure 4.5(c) as a preference ranking (along the y-axis) for three individuals along the political continuum from liberal (left side of the x-axis) to conservative (right side of the x-axis), for example, it is clear that one person prefers liberal candidates most and conservatives least, one person prefers conservative candidates most and liberals least, and one person prefers moderates most (with nearly equal disdain for conservatives and liberals). But they all agree on one thing: the moderate candidate is not the *worst*. Multi-peaked profiles, as in Figure 4.5(d) have no such locus of agreement.

Sen provided a way to capture this notion of disagreement through his "Condition of value restriction" (modified for my purposes in (162)).

(162) *Condition of value restriction*

In any triple (x, y, z) there is some alternative x such that all relevant voters agree

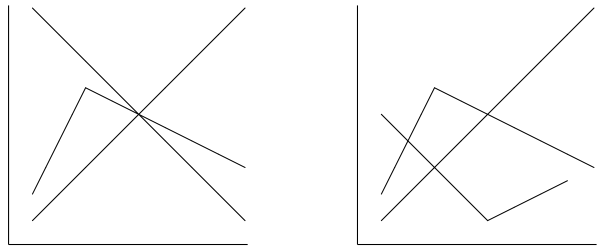

(c) Single-peaked rankings　(d) Multi-peaked rankings

Figure 4.5: Single- and multi-peaked rankings

that it is not worst, or agree that it is not best, or agree that it is not in the middle, i.e.

for all i: xP_iy or xP_iz or

for all i: yP_ix or zP_ix or

for all i: $[xP_iy$ and $xP_iz]$ or $[yP_ix$ and $zP_ix]$

Sen then showed that, if every triple of rankings satisfies the condition in (162), the aggregation function escapes Arrow's impossibility result.

The upshot is this: vague predicates, on the indecisional view, involve component dimensions that vary widely on their rankings. Where such rankings agree in substantial part (at least as much as reflected in (162)), it is possible to achieve an aggregated ranking without abandoning the assumption of transitivity. Therefore, where dimensions are apt to agree, or where standards are apt to agree, this account overgenerates vagueness effects.

The condition of value restriction is a restriction on the domain of alternatives, enforcing a certain uniformity among the voters (or dimensions, etc.) and violating Unrestricted Domain. This condition is therefore also related to the question of whether to adopt a "single-profile" or "multi-profile" version of Arrow's theorem. Arrow's theorem (particularly the Independence axiom) is framed in terms of the preferences voters could have rather than the preferences the voters actually have. His reason was epistemic: the social welfare function has to be chosen before it is known what people's preferences happen actually to be; if these could be anything, then it has to be ready for everything.[19] However, such a framing may not make sense for gradable adjectives: there is probably only one way for Reno and New York to be ranked in terms of any dimension relevant to the interpretation of *little*. This represents an implicit restriction on the domain of *little* from the viewpoint of Arrow's theorem.

One possible response in favor of the multi-profile version of Arrows theorem in the adjectival context is that we have good reason for gradable adjectives to be used consistently across contexts. Among other things, it facilitates the possibility of conversational coordination (both intra- and inter-speaker coordination). If contexts of use each specify a single profile, then we also have good reason to want coordination of the output among contexts. Of course, some criteria will not diverge on rankings across contexts (height, for example). But some will (level of interest, for example). Since the coordination of output among contexts is a job for the multi-profile version, we should prefer the multi-profile version.

Still, even accepting this motivation, it remains an open question whether there is enough variety among relevant profiles to obtain an Arrow-style impossibility result using interprofile constraints like Independence. In this case, there are single-profile impossibility results (relying on a intra-profile version of neutrality instead of Independence: suppose individual preferences for w vs. z are identical to those for x vs. y. Then the social preference for w vs. z must be identical to the social preference for $x$vs. y) that may apply. This avenue of investigation is left for future work.

19. Thanks to Michael Morreau for bringing this issue to my attention.

CHAPTER 5

LINGUISTICS AND LEGISLATIVE INTENT

The indecisional approach is rooted in an analogy to a legislature: as voters are to proposals, so dimensions are to entities. The legislature ranks proposals by aggregating voters' preferences among proposals. Gradable adjectives rank entities by aggregating dimensions' rankings of entities. The legislature's aggregation procedure is subject to seemingly reasonable constraints that produce paradox. Gradable adjectives' aggregation procedure is subject to seemingly reasonable constraints that produce paradox. In both cases, paradoxes are induced by intransitive aggregated rankings, despite the fact that each individual voter's preference, or each individual dimension, is transitive.

This chapter pursues the analogy in the other direction. If speaker's intent is relevant to the interpretation of vague gradable adjectives, so legislative intent is relevant to the interpretation of legislative commands. In other words, a judge interpreting a vague multidimensional adjective like *healthy* in a will authored by an individual and a judge interpreting a vague word like *reasonable* in a statute authored by the multiple individuals that comprise a legislature are engaged in the same fundamental tasks. If reference to speaker's intent is a licit interpretive tool in interpreting a will, then reference to legislative intent is a licit interpretive tool in interpreting a statute. This conclusion runs contrary to the dominant view in American law, a view whose consequences can be measured in (among other things) prison terms.

I begin by rehearsing the social choice argument against the use of legislative intent. In section 5.2, I show that the indecisional approach to vagueness bears upon the problem of identifying legislative intent. Consequently, in section 5.2.1, I defend courts' use of individual intent when interpreting vague predicates in single-authored texts like contracts, wills, and other documents. Many of these same defenses will apply to the use of legislative intent. Therefore, the social choice argument provides no good reason for disregarding legislative intent.

5.1 The argument against the use of legislative intent

For around 100 years, the dominant approach to interpreting statutes in American law was to give effect to the intent of the legislature (Manning, 2010; Hart & Sacks, 1958). Courts frequently affirmed, for example, that "the object of construction, as has been often said by the courts and writers of authority, is to ascertain the legislative intent, and, if possible, to effectuate the purposes of the lawmakers."[1]

By the early 1980s, the insights of social choice theory had begun to inspire a backlash against the use of legislative intent in statutory interpretation (Manning, 2010). In particular, two areas of economic research contributed to this backlash: interest group theory and social choice theory. The former undermined the soundness of relying on certain portions of legislative history (Scalia, 1998) by, for example, questioning legislators' practice of adding un-voted-on glosses to statutory text in official publications like congressional reports. Social choice theory, on the other hand, cast doubt on the very idea of a coherent legislative intent. The alleged incoherence of legislative intent reverberates throughout debates over its legal validity (Dworkin, 1986; Waldron, 1999; Ekins, 2012), and this incoherence almost always reduces to the following insight: there is no clear way to aggregate many individual intentions into one collective intent.[2]

The social choice argument is grounded on Arrow's Theorem (Arrow, 1950), a general-

1. *ICC v. Baird*, 194 U.S. 25, 38 (1904) (cited in Manning (2010)).

2. To be sure, there are arguments against the use of legislative intent that do not directly touch on the aggregation problem. One such argument focuses on the qualitative nature of mental states. Since mental states are subjective, conscious experiences, and since intention is a mental state (this argument presumes), it is not possible for a group to possess an intention (Searle, 1995). See List & Pettit (2011) for persuasive arguments against this view. Alternatively, Scalia (1998) has famously argued that democratic legitimacy (including bicameralism and presentment) attaches to statutes, not intentions. But the aggregation problem stalks the legitimacy argument, too: democratic legitimacy is satisfied (in part) by bicameralism, bicameralism is satisfied when both Houses of Congress pass a proposal by majority vote, and majority vote is subject to a version of the aggregation problem.

ization of Condorcet's paradox, repeated informally below for the legislative case (compare (63a)). Let $X = \{x, y, z, \ldots\}$ be candidates in a voting procedure (i.e., alternatives) and let $N = \{1, \ldots, n\}$ be the set of voters.

(163) a. Unanimity: If every voter prefers alternative x to y, the social preference prefers x to y.

b. Independence of Irrelevant Alternatives: The aggregated ranking of candidates $x \succcurlyeq y \succcurlyeq, \ldots, \succcurlyeq z$ should depend only the individuals' ranking of candidates $x \succcurlyeq_1 y \succcurlyeq_1, \ldots, \succcurlyeq_1 z$.

c. Unrestricted Domain: The voters can rank the relevant candidates in any way.

d. Non-Dictatorship: No one voter determines the result of the aggregation of preferences.

(164) Arrow's Theorem: There is no complete and transitive collective decision procedure that respects all four constraints.

In the legislative context, Arrow's conditions are normatively appealing. For example, Unanimity and Non-Dictatorship are desirable properties for a reasonably democratic voting procedure to have. Unrestricted Domain is also sensible, though it turns out that if you disallow voters from having certain preferences, you can turn the impossibility result into a possibility result (Black et al., 1958; Sen, 1970; Gaertner, 2001).

In contrast, the legislative version of Independence of Irrelevant Alternatives ("Independence") is more controversial. On the one hand, Independence intuitively cabins the aggregated ranking of candidates x and y to the "inherent" merits of x and y (without consideration of the "irrelevant" alternative z). The aggregation method "has to take each pair of alternatives separately, paying no attention to preferences for alternatives other than them" (Morreau, 2015). On the other hand, Independence requires us to throw away valuable information.[3] In particular, Independence's pairwise focus allows the aggregation function

3. See chapter 3.

to ignore the fact that voters themselves have transitive preferences (by hypothesis) (Saari, 2008).

As in the case of vagueness phenomena associated with gradable adjectives like *healthy*, Independence is in many ways at the heart of the social choice attack on legislative intent. Importantly, legislatures often abandon Independence in practice. One form of this abandonment is "agenda setters," who determine what gets voted on and when (committee chairs, for instance). Agenda setters may design the voting process such that their preferred alternative wins. As Shepsle (1992, 244) notes, in the Arrowian framework, "a clever agenda setter can produce any majority rule result she wants." But then which proposal wins depends not only on its position relative to the losing options; it also depends on the order in which proposals come to a vote. This is a violation of Independence, the assumption that a decision between two options should depend only on the voters' relative ranking of those two options (and not, for instance, on the order in which the options were considered).

In the legislative context, the problem with violations of Independence is that they make intent difficult to discern. Shepsle (1992, 248) argues that the tradeoff represented by agenda setting—trading away Independence for transitivity—"does not permit us to differentiate the 'will of the majority' from the machinations . . . of agenda setters." And Easterbrook (1983, 548) argues that violations of Independence "submerge" the legislative process, and "courts lose the information they need to divine the [legislature's] design."

Violations of Independence generally have this intent-obfuscating character. These violations invite strategic behavior (through the consideration of irrelevant alternatives, for example), and strategic behavior obstructs a court's ability to make inferences about the legislature's intent. After all, how is a court to reconstruct the collective preferences of the legislature when potentially hidden, seemingly irrelevant policy alternatives influence these preferences? For courts hunting legislative intent, violations of Independence create an inferential morass.

It follows, according the social choice argument, that the concept of legislative intent is meaningless. To avoid cycling, the legislative process gives up Independence. But giving up Independence has the effect of hopelessly muddying the waters in which courts search for legislative intent.

5.2 Vagueness and the social choice argument

The social choice argument against legislative intent is too strong. Turning back to the indecisional approach, it is clear that the social choice argument against legislative intent implies that vague multidimensional adjectives are meaningless. The argument against the use of legislative intent in statutory interpretation rests on a now-familiar aggregation problem: since there is no clear way to aggregate many legislators' individual intents into a unified intent, the legislature would appear to have no "legislative intent." This aggregation problem is formally analogous to the aggregation problem giving rise to adjectival vagueness.

As discussed in chapter 3, Arrow's Theorem is really about certain types of decisions: those that aggregate judgments of multiple alternatives along multiple criteria. Katz (2011, 7) explains, "It was understood from early on that Arrow's theorem is not just about voting but can be extended to tell us something far more general about all decision making, whether it involves collectivities or not."

Thus, the connection between vagueness and legislative intent therefore runs through Arrow's Theorem. Arrow's Theorem says that any decision-making procedure involving multiple criteria and multiple alternatives will have a problem like cycling, so long as certain assumptions are maintained. The interpretation of a vague predicate requires this type of decision (a choice, as discussed in chapter 3), and this decision involves multiple criteria (the relevant dimensions) and multiple alternatives (the members of the comparison class).

If the social choice argument against legislative intent rests on the same premises as the explanation for vagueness effects in multidimensional adjectives, then identifying the legislative intent behind a statutory command should be no harder than identifying the speaker's intent behind the use of a multidimensional adjective. In other words, on the indecisional view of vagueness phenomena, identifying legislative intent is no harder than interpreting vague words. This is a task speakers—including judges—perform every day. If courts can be trusted to interpret vague words like *healthy* in wills and contracts (and statutes), they should be trusted to identify the legislative intent behind a statutory provision. The social choice argument against the use of legislative intent in statutory interpretation fails.

5.2.1 Courts are trusted to interpret vague language

The law trusts courts to identify and interpret vague language. This is evident in the proliferation of vague legal standards like *reasonable*,[4] in legal doctrines that turn on the finding of "ambiguity" (often a cover term for vagueness),[5] and, especially, in the legal doctrine directly addressing vague language, the void-for-vagueness doctrine. This doctrine represents the principle that the legislature's use of "words and phrases [may be] so vague and indefinite that any penalty prescribed for their violation constitutes a denial of due process of law" under the Fifth and Fourteenth Amendments to the United States Constitution.[6]

The Supreme Court has observed three main problems with vague laws. First, vague laws fail to provide fair notice to the public as to what constitutes illegal conduct. The most famous recitation of this rationale for the void-for-vagueness doctrine is Justice Oliver Wendell Holmes, Jr.'s: "The law is full of instances where a man's fate depends on his estimating rightly, that is, as the jury subsequently estimates it, some matter of degree. If his judgment is wrong, not only may he incur a fine or a short imprisonment . . . he may

4. Generally, an individual is civilly liable for some wrong only if they have failed to act as a reasonable person in similar circumstances.

5. *See Hoffman Estates v. Flipside, Hoffman Estates*, 455 U.S. 489, 494 n.6 (1982) ("[T]he vagueness of a law affects overbreadth analysis. The Court has long recognized that ambiguous meanings cause citizens to 'steer far wider of the unlawful zone' . . . than if the boundaries of the forbidden areas were clearly marked."). In *Hoffman Estates*, the Supreme Court appears to use "ambiguous" and "vague" interchangeably.

6. *Champlin Ref. Co. v. Corporation Comm'n*, 286 U.S. 210, 243 (1932).

incur the penalty of death."[7] Second, vague laws represent an impermissible delegation of law-making authority to the judiciary. Therefore, vague laws challenge the separation of powers between the judicial and legislative branches. Third, and relatedly, vague laws fail to guide the discretion of executive officials and judges. In the Court's recent phrasing, "the Government violates [Due Process] by taking away someone's life, liberty, or property under a criminal law so vague that it fails to give ordinary people fair notice of the conduct it punishes, or so standardless that it invites arbitrary enforcement."[8]

Of course, the "so vague and indefinite" standard is itself vague, and its application requires that we trust courts to identify and interpret vague words. Many do not. For example, in his dissent in *Johnson v. United States*, Justice Thomas noted that the vagueness doctrine has been wielded by the Court "to achieve its own policy goals," from striking down economic regulations to striking down speech regulations to striking down abortion regulations.[9]

This distrust is easy to understand: so much legal language is vague, and so little of it is invalidated under the void-for-vagueness doctrine. The discrepancy undermines the claim that the void-for-vagueness doctrine is somehow about the "quantum" of vagueness involved (Post, 1994). In *Kolender v. Lawson*,[10] for instance, the Court held that a statute requiring a person to provide "credible" identification upon police request was void for vagueness. Yet in *Arizona v. Fulminante*, the Court articulated a similar standard for finding that a confession had been coerced in violation of the Fifth and Fourteenth Amendments: "Our cases have made clear that a finding of coercion need not depend upon actual violence by a government agent; a *credible* threat is sufficient."[11] In both cases, the law directs government actors to

7. *Nash v. United States*, 229 U.S. 373, 377 (1913).

8. *Johnson v. United States*, — U.S. — (2015).

9. For other criticisms along these lines, see Amsterdam (1968) and Goluboff (2010).

10. 461 U.S. 352 (1983).

11. 499 U.S. 279 (1991) (emphasis added).

make a determination of credibility (to assess certain kinds of identification in the first case and to avoid certain kinds of threats in the second). In both cases, an individual's potential incarceration hangs in the balance. But in one instance, the use of this multidimensional adjective is void for vagueness, and in the other instance, it is unquestionably part of the law.

However, the indecisional approach to vagueness elaborated in chapters 3 and 4 provides an explanation: the void-for-vagueness doctrine is best understood as policing the legality of the relevant dimensions involved in the interpretation of vague legal standards. In other words, the void-for-vagueness doctrine sorts those dimensions that are appropriate "triggers for police control" from those that are not (Post, 1994, 498).

In *Johnson v. United States*, for example, the Supreme Court confronted the Armed Career Criminals Act of 1984 ("ACCA"). Under this Act, a felon in possession of a firearm with three prior convictions for a "violent felony" qualifies for a mandatory fifteen-year sentence. The ACCA defines the term "violent felony" as a crime that is "burglary, arson, or extortion, involves use of explosives" or a crime that "otherwise involves conduct that presents a serious potential risk of physical injury to another." The Court struck down the "otherwise" clause (called the "residual" clause) on vagueness grounds.

In prior ACCA cases, the Court had attempted to factor the residual clause into something useful. For example, in one case, the Court concluded that "violent" meant "purposeful" and "aggressive" conduct that presented a risk of physical injury to another.[12] In another prior case, the Court had looked to the risk presented by the "ordinary case" of crime at hand.[13] Importantly, these interpretations all addressed the permissible dimensions of the vague multicriterial predicate *violent*: conduct violent with respect to purpose, conduct violent with respect to aggression, conduct violent in the ordinary case. Yet *Johnson* ultimately decided that this latter dimension—conduct violent in the ordinary case—was

12. *Begay v. United States*, 553 U.S. 137 (2008).

13. *Sykes v. United States*, 564 U. S. 1 (2011).

impermissible. As the *Johnson* Court explained, the problem with the residual clause was that it encouraged judges to make guesses about "violence" based on hypothetical conduct: "It ties the judicial assessment of risk to a judicially imagined 'ordinary case' of a crime." Concluding that hypothetical conduct was an impermissible dimension of "violence," the Court stuck down the residual clause as unconstitutionally vague.

The view from multidimensionality suggests that courts implicitly both recognize and exploit the connection between vagueness and multicriteriality in statutory interpretation.[14] Indeed, when confronted with a vague legal standard like *reasonable*, some courts respond with "factors tests," the cutting up of the meaning of the standard into its component criteria. For example, in certain actions, a statutory award of "reasonable" attorney's fees to the victorious party is mandatory. In order to determine whether fees are reasonable, the Ninth Circuit consults

> (1) the time and labor required, (2) the novelty and difficulty of the questions involved, (3) the skill requisite to perform the legal service properly, (4) the preclusion of other employment by the attorney due to acceptance of the case, (5) the customary fee, (6) whether the fee is fixed or contingent, (7) time limitations imposed by the client or the circumstances, (8) the amount involved and the results obtained, (9) the experience, reputation, and ability of the attorneys, (10) the undesirability of the case, (11) the nature and length of the professional relationship with the client, and (12) awards in similar cases.

Thus, when encountering a vague standard like *reasonable*, courts in the Ninth Circuit factor the meaning of this standard into its component dimensions.

---

14. Goluboff (2010) persuasively explains another void-for-vagueness case, *Papachristou v. Jacksonville*, 405 U.S. 156 (1972), along these lines. There, considering an anti-vagrancy statute, the Court found that "life style"—a dimension of vagrancy—impermissibly triggered police control.

5.2.2 *Intent and the interpretation of vague terms*

If courts are trusted to apply vague terms, one of the tools they use to ensure they apply vague terms correctly is individual intent. Individual intent has long been canonized as an interpretive tool. Courts interpret contracts "according to the intent of the parties,"[15] wills according to "the intention of the testator,"[16] and regulations promulgated by an agency head according to "the [agency head]'s intent at the time of the regulation's promulgation."[17] They sometimes interpret precedent according to the intentions of its specific author, even when that author is technically a multi-member court.[18] Thus, the social choice attack on intent applies to these practices, too, at least where this intent relates to the interpretation of vague predicates.

A bite-the-bullet response to this problem—that speaker's intent with respect to the use of certain gradable adjectives is meaningless—is unlikely to work. While linguists and philosophers debate whether speaker's intent is part of semantic or pragmatic meaning (or part of "narrow" or "broad" context) (see, e.g., Kaplan (1977); Bach (1999, 2001); King (2014)), there is little doubt that speaker's intent is an important factor in determining the meaning of contextually sensitive expressions and utterances.

First, speaker's intent is necessary to determine "what is said" by the use of demonstratives like *that woman* and anaphoric expressions like *she* and *him* (Grice, 1957; Kaplan, 1989).[19] For example, imagine a speaker pointing to a dog and saying "that dog is mean."

15. See, e.g., *Montana v. Wyoming*, 131 S. Ct. 1765, 1772 n.4 (2011). Contract formation involves the aggregation of multiple intents (at least those of the parties). Therefore, the social choice broadside against legislative intent should apply with special force to contract interpretation.

16. See, e.g., *Smith v. Bell*, 31 U.S. 68, 74 (1832).

17. See, e.g., *Thomas Jefferson Univ. v. Shalala*, 512 U.S. 504, 512, (1994).

18. See, e.g., *United States v. Belle*, 593 F.2d 487, 504 (3d Cir. 1979) ("But if Justice Brennan had intended to exclude linkage testimony from the protections of the Confrontation Clause, surely he would have chosen to restrict, or at least to distinguish, his holding in Douglas to the contrary.").

19. This discussion sets indexicals like *I* and *here* aside.

The speaker's pointing alone will not serve to fix the reference of *that dog*, because the speaker could be pointing slightly to the side of the dog, or at a part of the dog, or at a flea on the dog. In a case like this, it seems that the dog is fixed as the value of *that dog* precisely because the hearer has good reason to think that the speaker intended to talk about the dog. King (2014, 230) extends this reasoning to other contextually sensitive expressions "for which it is plausible that speaker intentions play some role in supplying them with values in context" including "tense, relational terms ('enemy', 'local'), *gradable adjectives*, quantifiers (domain restriction), modals, possessives ('John's book'), and so on."[20] In these cases, "it seems plausible that the value the expression takes in a context is the one that the speaker intends and that a competent, attentive, reasonable hearer who knows the common ground of the conversation would take the speaker to intend" (*id.*).

Second, intention-reading is especially important in imperatives, the language of the law. Imperatives are action-guiding with respect to future actions of the addressee. As Portner (2004) notes, "[I]mperatives represent actions which the addressee should take." In aid of this action-guiding quality, for example, the action must not already be complete (165).

(165) # The window is already open. Open the window!

In addition, imperatives display "functional heterogeneity," or the ability to be associated with a wide variety of speech acts (166) (Condoravdi & Lauer (2012) citing Schmerling (1982)).

(166)	a. (Mother to child) Go to bed	ORDER
	b. Have a cookie	INVITATION
	c. Go ahead	PERMISSION
	d. Take more vitamins	ADVICE
	e. . . .	. . .

20. Emphasis mine.

In Condoravdi & Lauer's account, imperatives express speaker preferences of a particular kind. First, the authors define a *preference structure* as a ranking of preferences over an information state. Formally, this looks something like (167).

(167) A preference structure relative to an information state W is a pair $\langle PS, \leq \rangle$ where $PS \subseteq \mathcal{P}(W)$ and $\leq$ is a weak partial order on PS

An individual may have more than one preference structure, corresponding to different types of preferences. In order to resolve potential conflicts among these different preferences, and thereby distinguish the preferences on which an agent should act, Condoravdi & Lauer introduce a distinguished, consistent preference structure for an agent. This preference structure is called the "effective preference structure," and it consolidates the agent's many preference structures into a consistent whole, as defined in (168) over the set of an agent's preference structures $\mathbb{P}$.[21]

(168) The effective preference structure for agent A is $\langle P_A, \leq_{P_A} \rangle$, where $P_A \subseteq \bigcup \mathbb{P}$

An imperative uttered by a speaker S is then defined as in (169).

(169) $[\![\text{IMP}(\phi)]\!]^c = \{w \mid \phi \text{ is a maximal element of } P_S \text{ at } w\}$

Thus, an imperative denotes those worlds where the speaker is committed to act as though ϕ is a maximal element of her effective preference structure.[22]

Functional heterogeneity falls out from contextual factors and assumptions about speaker and addressee desires. An expression of speaker preference is interpreted as an ORDER, for example, in some circumstances and as a request in others. As relevant here, then, speaker's intention is especially important in interpreting commands and directives incorporating gradable adjectives like (170), as well as in determining whether such commands are in fact commands at all.

21. I have omitted the operations intended to ensure consistency.

22. And in some contexts, the authors note, an imperative commits the speaker to act as if she prefers that the addressee prefer ϕ.

(170) Don't issue credible threats. ADVICE OR PROHIBITION

Third (and relatedly), reference to speaker's intent is necessary to determine what effect the speaker intended to achieve in making non-imperative utterances. Grice referred to this sort of intent as "communicative intent." According to Grice (1957, 220), "S meant something by *u*" is equivalent to "S intended the utterance of *u* to produce some effect in an audience by means of the recognition of this intention." In (171), for instance, Speaker B's response could be interpreted as a "yes" or a "no," depending on Speaker B's intended effect on Speaker A, and Speaker A's recognition of this intent. In other words, Speaker A must refer to Speaker B's intent in order to figure out what Speaker B means.

(171) Speaker A: Will you have some coffee?

Speaker B: Coffee would keep me awake.

Of course, communicative intent also affects the use of vague multidimensional adjectives. If someone describes a new book to you as "typical of its author," this may be a recommendation or a critique, depending on the effect the speaker intends her utterance to have.

It is therefore unlikely that we can excise individual intent from our interpretive practices, even in the small corner of the grammar devoted to multidimensional adjectives. Nor do advocates of the social choice critique of legislative intent want to abandon the concept of individual intent in the interpretation of contracts. For instance, Easterbrook (1983, 540) writes, "In choosing background rules for understanding and completing of contracts, courts ordinarily select the options they think the parties would have picked had they thought of the subsequently surfacing problems and been able to bargain about them beforehand at no cost." The social choice attack is intended to extend to statutory interpretation and no further.

However, the power of this attack exceeds its intended confines. And while a bite-the-bullet response might conclude that legislative intent and individual intent are equally suspicious and neither should be used, the necessity of individual intent in successful communication forecloses this response.

5.2.3 The usefulness of legislative intent

If recourse to individual intent is necessary for successful communication, and if we trust courts to look to individual intent to resolve the meaning of vague words in wills and the like, then we should trust courts to look to legislative intent when interpreting a vague statute.

First, the parallel between legal interpretation and ordinary communication is hard to avoid.[23] "[O]ur laws can be no more than efforts at communication based on the intention of the drafters" (Solan, 2010, 88). Evidence that acts of legislation are acts of communication comes from the "canons of construction," rules that guide statutory interpretation. Some canons of construction are designed to formalize common principles of English usage (Nelson, 2011). These canons include some of the familiar standbys of statutory interpretation, like "the ordinary meaning principle," "the presumption of consistent usage," and *expressio unius est exclusio alterius*, the idea that the legislature's express identification of members of a certain class is an implied rejection of other members of that class. Crucially, these canons are all attempts to attribute intent to the legislature.

The ordinary meaning principle, for example, provides that "words used in a statute are to be given the ordinary meaning in the absence of persuasive reasons to the contrary," and "the 'ordinary meaning' that matters is the common import of the words in American English at the time Congress enacted the statute" (Nelson, 2011, 83). This addresses the underdetermination of reference by meaning, a problem typically resolved by recourse to intent. For example, in *Gustafson v. Alloyd Co.*,[24] the Supreme Court decided the meaning of the word "prospectus" in the Securities Act of 1933. Crucially, the understanding of the word "prospectus" had changed since the 1933 Act's passage. In applying the ordinary meaning canon, the Court resolved this problem by looking to the understanding of "prospectus" at the time of the 1933 Act's enactment. The ordinary meaning canon therefore functions as a rule for attributing a particular type of communicative intent to the legislature, the

23. But see Dworkin (1986); Eskridge (1994); Scalia (1998).

24. 513 U.S. 561, 575-76 (1995) (discussed in Nelson (2011, 84)).

intent that the semantic content of "prospectus" should be determined at the time of enactment. Importantly, the Court reached this conclusion even though the 1933 Act contained a definitional section that defined "prospectus" in a way hostile the Courts interpretation.

The maxim *expressio unius est exclusio alterius* (the expression of one thing is the exclusion of another) also addresses a problem typically resolved by recourse to intent. This canon invites a negative inference from the statutory language. In *Leatherman v. Tarrant County Narcotics Intelligence and Coordination Unit*,[25] the Supreme Court held that Federal Rule of Civil Procedure 9(b)'s heightened requirements for "fraud or mistake" claims only applies to the types of claims listed. In particular, on the strength of the *expressio* canon, courts are without power to subject other claims to Rule 9(b)'s requirements. This is not at all different from Speaker B's response in the following dialogue.

(172) Speaker A: I like John and Sonia.
Speaker B: I like John.

The clear inference—based on the effect Speaker B likely intends to produce in her audience– is that Speaker B does not like Sonia.

The canons constitute rules for attributing communicative intent to the legislature. And such rules are only necessary if statutes are treated as efforts at communication.[26]

In addition, legal interpreters frequently treat groups of individuals—like the legislature— as having the communicative purposes and goals of individuals. Locutions like "Congress intends" or "the legislature meant" appear in American legal opinions at a frequency of about 60,000 times per decade for the last two decades (Solan, 2010, 101). Even those normally hostile to reliance on legislative intent use these locutions. The reason, Solan concludes (p. 104), is that "it is almost impossible to avoid thinking in . . . intentionalist terms." The seemingly unproblematic attribution of intents to collectives in everyday communication also

25. 507 U.S. 163, 168 (1993) (discussed in Nelson (2011, 89)).

26. It is also no accident that the stable of canons is flexible enough to arrive at contradictory attributions of communicative intent (Llewellyn, 1949).

suggests that an attribution of intent to the legislature is similarly unproblematic.

Finally, without some notion of legislative intent, it is impossible to make sense of legislative errors. In *United States v. Granderson*,[27] for example, the defendant was sentenced to five years probation, though he faced up to six months in prison. When he violated the terms of his probation, he was resentenced under a statute that provided the following: "If the defendant violates a condition of probation ... the court shall revoke the sentence of probation and sentence the defendant to not less than one-third of the original sentence." The most natural reading of this statute is that Granderson should have received one-third of the original five-years-probation sentence—a more lenient sentence for violating the terms of his probation. The Court rejected this outcome as absurdly contrary to the legislature's intended meaning. This move is unavailable if legislative intent is not a legitimate tool of statutory interpretation.

Therefore, statutes at least approximate a form of communication: they communicate the legislature's intended meaning to the community (Ekins, 2012). And communication requires an attribution of communicative intent to the speaker. If the social choice critique is right, though, this attribution (whether to an individual or a collective speaker) is incoherent. According to Shepsle's argument, it is unjustifiable to infer from a speaker's use of a vague multidimensional adjective like *healthy* how the speaker would use that predicate to describe anything else. Yet such inferences are both communicatively essential and legally commonplace. The social choice critique cannot mean that the notion of individual intent is incoherent.[28]

27. 511 U.S. 39 (1994) (discussed in Solan (2010, 104–05)).

28. There is yet a small refuge for the legislative-intent skeptic. Perhaps it is possible to isolate communicative intent from other types of speaker intent, such that attributions of communicative intent to the legislature are coherent but attributions of other intents (intents to approach a public policy problem in a certain way, for instance) are incoherent. I am doubtful that this refuge is secure. First, the social choice critique provides no room for distinguishing between types of speaker intent. Second, this approach more resembles a careful, pragmatic approach to legislative intent, in the vein of McNollgast (1994) and Nourse

If the social choice critique fails in the case of individual intent, then it holds no special sway over the case of legislative intent. From the standpoint of social choice, there is no difference between "voters" and "dimensions," or between "proposals" and "members of the comparison class." Thus, there is no reason for the social choice attack to succeed against legislative intent but not against individual intent.

Of course, the collective nature of the legislative intent may merit some special consideration. When legislators' preferences for a proposal tend to fall along one liberal-conservative dimension, as Farber & Frickey (1988) discuss, legislative intent is more readily discernable. Similarly, McNollgast (1994) and Nourse (2012) concentrate on the importance that agenda setters may play in determining the intent of the legislature as a whole. However, the desire to attribute intent to the legislature—as reflected in canons of construction, common legal interpretive practice, and the possibility of legislative error—parallels the necessity of attributing intent to an individual speaker. In the face of these reasons for attributing intent to the legislature, the argument from social choice against the use of legislative intent in statutory interpretation must fail.

5.3 Legislative intent endures

The indecisional approach to vagueness begins with an analogy to the legislature. And just as the workings of a legislature inform the indecisional approach to vagueness phenomena, the indecisional approach informs how we should understand legislatures. If legislative intent is incoherent because of the aggregation problem, and if the indecisional approach to vagueness is correct, then speaker's intent must be incoherent as an interpretive tool for understanding the meaning of vague predicates. Since speaker's intent obviously—and unavoidably—informs the interpretation of these predicates, then legislative intent is not incoherent (or at least not because of the aggregation problem).

(2012). But in that case, the legislative-intent skeptic has become a legislative-intent believer.

The structure of this argument is a kind of *reductio ad absurdum*: the social choice critique can be used against the notion of individual intent; individual intent is not meaningless; therefore, this argument against individual intent must be wrong. But since there is no principled way to distinguish the argument against individual intent from the argument against legislative intent (at least according to social choice formalisms), and since we have independent reasons to treat legislative intent as sensible, the argument against legislative intent must be wrong, too. In other words, the social choice problems encountered in interpreting a statute are no greater than those encountered in interpreting a vague word in a will; if author's intent is a licit interpretive tool for one, it is licit for the other.

A positive account of legislative intent might infer joint intent from the structural features of joint action (List & Pettit, 2011; Ekins, 2012; Nourse, 2012). Such accounts place emphasis on the institutional features of the legislature—agenda setting, deliberation—as solutions to the inferential morass created by violations of Independence. For example, sequential procedures induce legislators to change their preferences in some cases, eliminating the aggregation problem.

The problem with these approaches, from the point of view developed here, is that it has no obvious analogue to the case of the individual. Perhaps sequential procedures—violations of Independence—resolve the problem of collective irrationality such that legislative intent is a meaningful guide in statutory interpretation. Even then, however, this fix is unlikely to be available to the problem to individual intent, and individual intent is not (as far as I know) subject to the same sort of structure-induced equilibrium.

But whatever the exact nature of legislative intent, it is not meaningless. And since the charge of meaninglessness undergirds, implicitly or explicitly, most prominent critiques of legislative intent, these critiques are faulty. A legislative command (usually) can and should be treated as one would treat the speech of an individual, at least when the individual's speech contains vague predicates. This follows not because legislatures are simple, but because, as the indecisional approach tells us, meaning is complex.

CHAPTER 6

CONCLUSION

The at-times pointy formalisms of the foregoing chapters do not tell the full story of vagueness, or honor its complexity, at least in my telling.

We know that a 5-foot tall man is not *tall* and that a 9-foot tall man is *tall*. At some point on the scale of height, we have traversed the distance between false and true, and yet we have indications from the apparent transitivity of indifference that the path between 0 and 1 is traversable without any big jumps (173).

(173) Any man 1mm taller than a tall man is tall.

Both cannot be right. Either there is no big jump between 0 and 1, meaning that the scale is continuous, or the scale is not traversable, meaning that indifference is in fact intransitive.

Furthermore, if we agree that the scale for the positive use of *tall* is the same as the scale for the comparative *taller*, the space of possible analyses is dramatically reduced. As discussed in chapters 2 and 3, theories of vagueness tend to agree on a choice rule for vague predicates R (174) that has the properties of a transitive strict relation $\succ$ and an intransitive indifference relation $\sim$.

(174) $[\![\text{happy}]\!]^c = \{x \mid \forall y \in c,\ xRy\}$

This strict relation preserves comparative semantics, and the intransitive indifference relation preserves room for vagueness effects.

Since Arrow's theorem involves assumptions of both transitivity *and* completeness, however, it may be the case that completeness fails rather than transitivity when other assumptions (Unrestricted Domain, Unanimity, Independence, Non-Dictatorship) hold. Indeed, certain philosophical theories tend to mimic this effect. As intimated in chapter 3, contextualism may be understood as denying completeness (at least in some sense) because certain relations between entities are excluded from consideration in context (Fara, 2002, 59). However, such an approach is difficult to reconcile with traditional semantics for gradable adjectives like *tall*. There is ample evidence (at least in my view) that the semantics of *tall* involves intransitive relations, but little evidence that two entities with some height may fail to be related by some tallness relation.

Additionally, the intransitivity of the indifference relation is sensible when the relata are multidimensional bundles, like vectors $(5, 3)$. As explained in chapter 2, a ranking of such bundles may include $(10, 2) \succ (9, 2)$, $(10, 2) \sim (5, 3)$, and $(9, 2) \sim (5, 3)$. Both $(9, 2)$ and $(10, 2)$ are indifferent to $(5, 3)$, but $(10, 2) \succ (9, 2)$. One of the latter two bundles is strictly preferred to the other, but both are indifferent to a third. This provides a sensible understanding of what a system would look like with a strict comparative and and intransitive indifference relation.

Yet while some semantic requirements drive the collective choice approach, this approach is also in tension with other common semantic assumptions. For example, the collective choice approach does not entail the existence of totally ordered scales of degrees of, e.g., healthiness typically invoked in the semantics of gradable adjectives like *healthy*(Cresswell, 1976b; von Stechow, 1984; Heim, 1985, 2001; Kennedy, 1999). Indeed, to the extent the collective choice approach relies on ordinal rankings of entities, these rankings (and not degrees) are conceptually primitive (Sassoon, 2013b, 80). Degrees are equivalence classes under the weak order $\succsim$, where an equivalence class is the set of entities related by a reflexive, transitive, and symmetric equivalence relation (for example, a degree of *height* is an equivalence class under the equivalence relation *equally tall*). However, it is perhaps most accurate to say that the social choice approach is agnostic with respect to degrees: in the voting context, it is possible to recreate Arrow's result in the context of cardinal utilities where interpersonal utilities are incomparable. Similarly, it is possible to recreate Arrow's result in the semantic context where the relevant dimensions are incommensurable, as properties like cholesterol and blood pressure (for *healthy*) generally are.

On the indecisional approach advanced here, multidimensionality is key to understanding vagueness phenomena. The question "how can continuity fail in a multidimensional space?" is very close in my view to the question "what is vagueness?" One way continuity can fail is when there is a hole in the space, as demonstrated by Chichilnisky (1982). And Chichilnisky also showed that that properties giving rise to holes are very similar to reasonable constraints on the semantics of gradable adjectives. Vague gradable adjectives are vague because multidimensionality plus these reasonable constraints leads to a discontinuity. If continuity is or is very similar to tolerance, then tolerance also fails in a multidimensional choice.

This account therefore vindicates Lewis's original observation. Social choice theory makes clear the sense in which vagueness effects result from paradoxes of decision making. The sorites paradox and Condorcet's paradox derive from the same source: it is impossible to make a collective decision over a multidimensional choice space under reasonable constraints like unanimity, anonymity, and continuity. So we waffle; we flip-flop; we vacillate. In other words, vagueness is semantic indecision.

REFERENCES

Alxatib, Sam, Peter Pagin & Uli Sauerland. 2013. Acceptable contradictions: Pragmatics or semantics? a reply to cobreros et al. *Journal of philosophical logic* 42(4). 619–634.

Alxatib, Sam & Francis Jeffry Pelletier. 2011. The psychology of vagueness: Borderline cases and contradictions. *Mind & Language* 26(3). 287–326. doi:10.1111/j.1468-0017.2011.01419.x.

Amsterdam, Anthony G. 1968. Note: The void for vagueness doctrine in the supreme court. *U. Pa. L. Rev.* 109. 67.

Arrow, Kenneth J. 1950. A difficulty in the concept of social welfare. *The Journal of Political Economy* 328–346. doi:10.1086/256963.

Arrow, Kenneth J. 1959. Rational choice functions and orderings. *Economica* 121–127.

Arrow, Kenneth J & Hervé Raynaud. 1986. *Social choice and multicriterion decision-making*. MIT Press.

Arrow, K.J. 1963. *Social choice and individual values*. Yale University Press 2nd edn.

Bach, Kent. 1999. The semantics-pragmatics distinction: What it is and why it matters. In Ken Turner (ed.), *The semantics-pragmatics interface from different points of view*, 65–84. Elsevier.

Bach, Kent. 2001. You don't say? *Synthese* 128(1-2). 15–44.

Baigent, Nicholas. 2010. Topological theories of social choice. In Kenneth J Arrow, Amartya Sen & Kotaro Suzumura (eds.), *Handbook of social choice and welfare*, vol. 2, 301–334. Elsevier.

Bale, Alan Clinton. 2008. A universal scale of comparison. *Linguistics and Philosophy* 31(1). 1–55.

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

Baryshnikov, Yuliy M. 1997. Topological and discrete social choice: in a search of a theory. In *Topological social choice*, 53–63. Springer. doi:10.1007/978-3-642-60891-9$_3$.

van Benthem, Johan. 1982. Later than late: On the logical origin of the temporal order. *Pacific Philosophical Quarterly anc Personalist (The) Los Angeles, Cal.* 63(2). 193–203.

Black, Duncan, Robert Albert Newing, Iain McLean, Alistair McMillan & Burt L Monroe. 1958. *The theory of committees and elections*. Springer.

Blackorby, Charles, David Donaldson & John A Weymark. 1984. Social choice with interpersonal utility comparisons: a diagrammatic introduction. *International Economic Review* 327–356.

Burnett, Heather. 2014. A delineation solution to the puzzles of absolute adjectives. *Linguistics and Philosophy* 37(1). 1–39. doi:10.1007/s10988-014-9145-9.

Chichilnisky, G. 1982. Social aggregation rules and continuity. *The Quarterly Journal of Economics* 97(2). 337–352. doi:10.2307/1880762.

Chichilnisky, Graciela. 1980. Social choice and the topology of spaces of preferences. *Advances in Mathematics* 37(2). 165–176. doi:10.1016/0001-8708(80)90032-8.

Chichilnisky, Graciela & Geoffrey Heal. 1983. Necessary and sufficient conditions for a resolution of the social choice paradox. *Journal of Economic Theory* 31(1). 68–87. doi:10.1016/0022-0531(83)90021-2.

Cobreros, Pablo, Paul Egré, David Ripley & Robert van Rooij. 2012. Tolerant, classical, strict. *Journal of Philosophical Logic* 41(2). 347–385. doi:10.1007/s10992-010-9165-z.

Cobreros, Pablo, Paul Egré, David Ripley & Robert van Rooij. 2015. Vagueness, truth and permissive consequence. In *Unifying the philosophy of truth*, vol. 36 Logic, Epistemology, Logic, Epistemology, and the Unity of Science, 409–430. Springer. doi:10.1007/978-94-017-9673-6$_2$1.

Condoravdi, C. 2009. Measurement and intensionality in the semantics of the progressive .

Condoravdi, Cleo & Sven Lauer. 2012. Imperatives: Meaning and illocutionary force. *Empirical issues in syntax and semantics* 9. 37–58.

Cresswell, Max J. 1976a. The semantics of degree. In *Montague grammar*, 261–292. Academic Press, New York.

Cresswell, Max J. 1976b. The semantics of degree. *Montague grammar* 261–292.

Douven, Igor, Lieven Decock, Richard Dietz & Paul Égré. 2013. Vagueness: A conceptual spaces approach. *Journal of Philosophical Logic* 42(1). 137–160. doi:10.1007/s10992-011-9216-0.

Dowty, D.R. 1979. *Word meaning and montague grammar: The semantics of verbs and times in generative semantics and in montague's ptq*, vol. 7. Springer. doi:http://dx.doi.org/10.1007/978-94-009-9473-7.

Dummett, Michael. 1975. Wang's paradox. *Synthese* 30(3-4). 301–324.

Dworkin, Ronald. 1986. *Law's empire*. Harvard University Press.

Easterbrook, Frank H. 1983. Statutes' domains. *The University of Chicago Law Review* 533–552.

Égré, Paul & Nathan W Klinedinst. 2011. Introduction. In Paul Égré & Nathan W Klinedinst (eds.), *Vagueness and language use*, Palgrave Macmillan.

Ekins, Richard. 2012. *The nature of legislative intent*. Oxford University Press.

Eskridge, William N. 1994. *Dynamic statutory interpretation*. Harvard University Press.

Fara, Delia Graff. 2002. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics* 28(1). 45–82. doi:10.5840/philtopics20002816. Originally published under the name Delia Graff.

Fara, Delia Graff. 2008. Profiling interest relativity. *Analysis* 68(300). 326–335.

Farber, Daniel A & Philip P Frickey. 1988. Legislative intent and public choice. *Virginia Law Review* 423–469.

Fine, K. 1975. Vagueness, truth and logic. *Synthese* 30(3). 265–300. doi:10.1007/BF00485047.

Gaertner, Wulf. 2001. *Domain conditions in social choice theory*. Cambridge Univ Press.

Gaertner, Wulf. 2009. *A primer in social choice theory*. Oxford University Press.

Gärdenfors, Peter. 2004. *Conceptual spaces: The geometry of thought*. MIT Press.

Giannakidou, Anastasia. 2004. Domain restriction and the arguments of quantificational determiners. In Robert B. Young (ed.), *Semantics and linguistic theory*, vol. 14, 110–126. doi:10.3765/salt.v14i0.2902.

Goluboff, Risa L. 2010. Dispatches from the supreme court archives: Vagrancy, abortion, and what the links between them reveal about the history of fundamental rights. *Stan. L. Rev.* .

Greene, Joshua D, R Brian Sommerville, Leigh E Nystrom, John M Darley & Jonathan D Cohen. 2001. An fmri investigation of emotional engagement in moral judgment. *Science* 293(5537). 2105–2108. doi:10.1126/science.1062872.

Grice, H Paul. 1957. Meaning. *The philosophical review* 377–388.

Grinsell, Timothy W. 2014. Unlikely imperfectives. In Todd Snider, Sarah D'Antonio & Mia Weigand (eds.), *Semantics and linguistic theory*, vol. 24, 395–413. doi:10.3765/salt.v24i0.2431.

Grinsell, Timothy Wood. 2013. Avoiding predicate whiplash: social choice theory and linguistic vagueness. In Anca Chereches (ed.), *Semantics and linguistic theory*, vol. 22, 424–440. doi:10.3765/salt.v22i0.2628.

Hart, Henry Melvin & Albert Martin Sacks. 1958. *The legal process: Basic problems in the making and application of law*, vol. 2.

Heal, Geoffrey. 1997. Social choice and resource allocation: a topological perspective. In Geoffrey Heal (ed.), *Topological social choice*, 1–14. Springer. doi:10.1007/978-3-642-60891-9_1.

Heim, Irene. 1985. Notes on comparatives and related matters. *Unpublished ms., University of Texas, Austin* .

Heim, Irene. 2001. Degree operators and scope. *Audiatur vox sapientiae. a festschrift for Arnim von Stechow* 214–239.

Hoeksema, Jacob. 1995. The semantics of exception phrases. *Quantifiers, Logic, and Language* 145–177.

Hyde, Dominic. 2008. *Vagueness, logic and ontology.* Ashgate Publishing, Ltd.

Janich, Klaus. 1984. *Topology.* Springer Verlag.

Kahneman, Daniel & Amos Tversky. 1984. Choices, values, and frames. *American psychologist* 39(4). 341.

Kamp, Hans. 1975. Two theories about adjectives. In Edward Keenan (ed.), *Formal semantics of natural language*, 123–155. Cambridge University Press. doi:10.1017/CBO9780511897696.011.

Kamp, Hans. 1981. A theory of truth and semantic representation. In Paul Portner & Barbara H. Partee (eds.), *Formal semantics: the essential readings*, 189–222. Blackwell. doi:10.1002/9780470758335.

Kamp, Hans & Barbara Partee. 1995. Prototype theory and compositionality. *Cognition* 57(2). 129–191. doi:10.1016/0010-0277(94)00659-9.

Kaplan, David. 1977. Demonstratives: An essay on the semantics, logic, metaphysics, epistemology of demonstratives and other indexicals. In J. Almog, J. Perry & H. Wettstein (eds.), *Themes from kaplan*, Oxford University Press.

Kaplan, David. 1989. Afterthoughts. In J. Almog, J. Perry & H. Wettstein (eds.), *Themes from kaplan*, Oxford University Press.

Katz, L. 2011. *Why the law is so perverse.* University Of Chicago Press.

Keefe, R. 2000. Theories of vagueness. 2000. *Cambridge Studies in Philosophy* .

Kennedy, Christopher. 1999. *Projecting the adjective: The syntax and semantics of gradability and comparison.* Routledge.

Kennedy, Christopher. 2001. Polar opposition and the ontology of degrees. *Linguistics and philosophy* 24(1). 33–70.

Kennedy, Christopher. 2007. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy* 30(1). 1–45. doi:10.1007/s10988-006-9008-0.

Kennedy, Christopher. 2011. Vagueness and comparison. *Vagueness and language use* 73–97.

Kennedy, Christopher & Louise McNally. 2005. Scale structure, degree modification, and the semantics of gradable predicates. *Language* 345–381. doi:10.1353/lan.2005.0071.

King, Jeffrey C. 2014. Speaker intentions in context. *Noûs* 48(2). 219–237.

Klein, Ewan. 1980. A semantics for positive and comparative adjectives. *Linguistics and philosophy* 4(1). 1–45. doi:10.1007/BF00351812.

Kratzer, A. 1981. The notional category of modality. In Hans J. Eikmeyer & Hannes Rieser (eds.), *Words, worlds, and contexts*, 38–74. De Gruyter. doi:10.1515/9783110842524-004.

Lakoff, George. 1970. Global rules. *Language* 627–639.

Landman, F. 1992. The progressive. *Natural language semantics* 1(1). 1–32. doi: http://dx.doi.org/10.1007/BF02342615.

Lassiter, Daniel & Noah D Goodman. 2015. Adjectival vagueness in a bayesian model of interpretation. *Synthese* 1–36. doi:10.1007/s11229-015-0786-1.

Lauwers, Luc. 2000. Topological social choice. *Mathematical social sciences* 40(1). 1–39. doi:10.1016/S0165-4896(99)00036-0.

Lauwers, Luc. 2009. The topological approach to the aggregation of preferences. *Social Choice and Welfare* 33(3). 449–476. doi:10.1007/s00355-009-0373-9.

Lewis, David. 1970. General semantics. *Synthese* 22(1). 18–67.

Lewis, David. 1986. *On the plurality of worlds*. Oxford: Blackwell.

List, Christian & Philip Pettit. 2011. *Group agency: The possibility, design, and status of corporate agents*. Oxford University Press Oxford.

Livnat, Adi & Nicholas Pippenger. 2006. An optimal brain can be composed of conflicting agents. *Proceedings of the National Academy of Sciences of the United States of America* 103(9). 3198–3202. doi:10.1073/pnas.0510932103.

Llewellyn, Karl N. 1949. Remarks on the theory of appellate decision and the rules or canons about how statutes are to be construed. *Vand. L. Rev.* 3. 395.

Luce, R Duncan. 1956. Semi-orders and a theory of utility discrimination. *Econometrica, Journal of the Econometric Society* 178–191. doi:10.2307/1905751.

Manning, John F. 2010. Second-generation textualism. *Cal. L. Rev.* 98. 1287.

Mas-Colell, Andreu, Michael Dennis Whinston, Jerry R Green et al. 1995. *Microeconomic theory*, vol. 1. Oxford University Press.

McNally, L. 2011. The relative role of property type and scale structure in explaining the behavior of gradable adjectives. In Rick Nouwen, Robert van Rooij, Uli Sauerland & Hans-Christian Schmitz (eds.), *Vagueness in communication*, 151–168. Springer. doi: 10.1007/978-3-642-18446-8_9.

McNollgast. 1994. Legislative intent: the use of positive political theory in statutory interpretation. *Law and Contemporary Problems* .

Mehta, Paras. 1997. Topological methods in social choice: an overview. In Geoffrey Heal (ed.), *Topological social choice*, 87–97. Springer. doi:10.1007/978-3-642-60891-9_5.

Moltmann, Friederike. 2005. Comparatives without degrees: a new approach. In *Proceedings of the fifteenth amsterdam colloquium*, 155–160.

Montague, Richard. 1974. Formal philosophy; selected papers of richard montague .

Morreau, Michael. 2015. Theory choice and social choice: Kuhn vindicated. *Mind* 124(493). 239–262. doi:10.1093/mind/fzu176.

Nelson, Caleb. 2011. *Statutory interpretation*. Foundation Press Thomson/West.

Nourse, Victoria. 2012. A decision theory of statutory interpretation: Legislative history by the rules. *Yale Law Journal* 122. 70–152.

Okasha, Samir. 2011. Theory choice and social choice: Kuhn versus arrow. *Mind* 83–115. doi:10.1093/mind/fzr010.

Portner, P. 1998. The progressive in modal semantics. *Language* 760–787. doi: http://dx.doi.org/10.2307/417002.

Portner, Paul. 2004. The semantics of imperatives within a theory of clause types. In *Proceedings of semantics and linguistic theory*, vol. 14, 235–252. CLC Publications Ithaca, NY.

Post, Robert C. 1994. Reconceptualizing vagueness: Legal rules and social orders. *California Law Review* 491–507.

Potts, Christopher. 2005. *The logic of conventional implicatures*. Oxford University Press.

Qing, Ciyang & Michael Franke. 2014. Gradable adjectives, vagueness, and optimal language use: A speaker-oriented model. In Todd Snider, Sarah D'Antonio & Mia Weigand (eds.), *Semantics and linguistic theory*, vol. 24, 23–41. doi:10.3765/salt.

Raffman, Diana. 1994. Vagueness without paradox. *The Philosophical Review* 103(1). 41–74. doi:10.2307/2185872.

Raffman, Diana. 1996. Vagueness and context-relativity. *Philosophical Studies* 81(2). 175–192.

Reich, Ingo. 2004. Association with focus and choice functions–a binding approach. *Research on Language and Computation* 2(3). 463–489. doi:0.1007/s11168-004-0902-8.

Reinhart, Tanya. 1997. Quantifier scope: How labor is divided between qr and choice functions. *Linguistics and philosophy* 20(4). 335–397. doi:10.1023/A:1005349801431.

Ripley, David. 2011. Contradictions at the borders. In Rick Nouwen, Robert van Rooij, Uli Sauerland & Hans-Christian Schmitz (eds.), *Vagueness in communication*, 169–188. Springer. doi:10.1007/978-3-642-18446-8$_1$0.

Rizza, Davide. 2013. Deconstructing a topological sorites. *Philosophia Mathematica* 21(3). 361–364. doi:10.1093/philmat/nkt025.

Roberts, FS. 1979. *Measurement theory*. Addison-Wesley, Reading, MA.

van Rooij, R. 2011. Vagueness and linguistics. In Giuseppina Ronzitti (ed.), *Vagueness: A guide*, 123–170. Springer. doi:10.1007/978-94-007-0375-9.

Rotstein, Carmen & Yoad Winter. 2004. Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics* 12(3). 259–288. doi: 10.1023/B:NALS.0000034517.56898.9a.

Russell, Bertrand. 1905. On denoting. *Mind* 479–493.

Saari, Donald G. 1997. Informational geometry of social choice. In Geoffrey Heal (ed.), *Topological social choice*, 65–86. Springer. doi:10.1007/978-3-642-60891-9$_4$.

Saari, Donald G. 2008. *Disposing dictators, demystifying voting paradoxes*. Cambridge University Press.

Sassoon, Galit W. 2013a. A typology of multidimensional adjectives. *Journal of semantics* 30(3). 335–380. doi:10.1093/jos/ffs012.

Sassoon, Galit W. 2013b. *Vagueness, gradability and typicality: The interpretation of adjectives and nouns*. Brill. doi:10.1163/9789004248588.

Sassoon, Galit W. 2015. A degree-approach account of multidimensional gradability. Available at semanticsarchive.net.

Scalia, Antonin. 1998. *A matter of interpretation: Federal courts and the law: Federal courts and the law*. Princeton University Press.

Schmerling, Susan. 1982. How imperatives are special, and how they aren't. *Papers from the Parasession on Nondeclaratives, Chicago Linguistic Society* 202. 218.

Schmidt, Lauren A., Noah D. Goodman, David Barner & Joshua B. Tenenbaum. 2009. How tall is tall? compositionality, statistics, and gradable adjectives. In *Proceedings of the thirty-first annual conference of the Cognitive Science Society*, 2759–2764.

Schwarzschild, Roger. 2008. The semantics of comparatives and other degree constructions. *Language and Linguistics Compass* 2(2). 308–331.

Searle, John R. 1995. *The construction of social reality*. Simon and Schuster.

Sen, Amartya. 1977. Social choice theory: A re-examination. *Econometrica: Journal of the Econometric Society* 53–89.

Sen, Amartya K. 1970. *Collective choice and social welfare*. North-Holland Publishing Co.

Shapiro, Stewart. 2006. *Vagueness in context*. Clarendon Press.

Shepsle, Kenneth A. 1992. Congress is a "they," not an "it": Legislative intent as oxymoron. *International Review of Law and Economics* 12(2). 239–256.

Smith, Nicholas JJ. 2008. *Vagueness and degrees of truth*. Oxford University Press.

Solan, Lawrence M. 2010. *The language of judges*. University of Chicago Press.

Solt, Stephanie & Nicole Gotzner. 2012. Experimenting with degree. In Anca Chereches (ed.), *Semantics and linguistic theory*, vol. 22, 166–187. doi:10.3765/salt.

Sorensen, Roy. 2001. *Vagueness and contradiction*. Oxford University Press.

Stanley, Jason. 2003. Context, interest relativity and the sorites. *Analysis* 63(280). 269–281.

von Stechow, Arnim. 1984. Comparing semantic theories of comparison. *Journal of semantics* 3(1). 1–77.

Stegenga, Jacob. 2013. An impossibility theorem for amalgamating evidence. *Synthese* 190(12). 2391–2411. doi:10.1007/s11229-011-9973-x.

Syrett, K., C. Kennedy & J. Lidz. 2010. Meaning and context in children's understanding of gradable adjectives. *Journal of semantics* 27(1). 1–35. doi: http://dx.doi.org/10.1093/jos/ffp011.

Toledo, Assaf & Galit W Sassoon. 2011. Absolute vs. relative adjectives-variance within vs. between individuals. In Neil Ashton, Anca Chereches & David Lutz (eds.), *Semantics and linguistic theory*, vol. 21, 135–154. doi:10.3765/salt.

Van Rooij, Robert. 2011. Implicit versus explicit comparatives. *Vagueness and language use* 51–72.

Venneman, Theo & Renate Bartsch. 1972. Semantic structures. *A study in the relation* .

Waldron, Jeremy. 1999. *Law and disagreement*. Oxford University Press.

Weber, Zach & Mark Colyvan. 2010. A topological sorites. *The Journal of Philosophy* 107(6). 311–325. doi:10.5840/jphil2010107624.

Weyl, E Glen. 2009. Whose rights? a critique of individual agency as the basis of rights. *Politics, Philosophy & Economics* 8(2). 139–171. doi:10.1177/1470594X09102235.

Williams, Robert. 2012. Vagueness. In Delia Graff Fara & Gillian Russell (eds.), *The routledge companion to philosophy of language*, Routledge.

Williamson, Timothy. 1994. *Vagueness*. Routledge.

Winter, Yoad. 1997. Choice functions and the scopal semantics of indefinites. *Linguistics and philosophy* 20(4). 399–467. doi:10.1023/A:1005354323136.

Zadeh, Lotfi A. 1975. Fuzzy logic and approximate reasoning. *Synthese* 30(3-4). 407–428. doi:10.1007/BF00485052.