Multidimensional Adjectives

JUSTIN D’AMBROSIO
University of St. Andrews

BRIAN HEDDEN
Australian National University

Multidimensional adjectives are ubiquitous in natural language. An adjective F is multidimensional just in case whether F applies to an object or pair of objects depends on how those objects stand with respect to multiple underlying dimensions of F-ness. Developing a semantics for multidimensional adjectives requires us to address *the problem of dimensional aggregation*: how do the application conditions of an adjective F in its positive and comparative forms depend on its underlying dimensions? Here we develop a semantics for multidimensional adjectives that incorporates aggregation functions. We then explore an analogy between dimensional aggregation and preference aggregation, bringing results from social choice theory to bear on the number and kind of aggregation functions that are admissible in a context. These results suggest that, for any given adjective, there will often be multiple aggregation functions admissible, meaning that multidimensional comparatives are often vague.

Keywords: gradable adjectives, multidimensionality, aggregation, vagueness, social choice theory

1. Introduction

Suppose I tell you that Suzy is healthy, or that she is healthier than Bill. What do I mean? What are the truth-conditions of my claims? We know that ‘healthy’ is a gradable adjective, since it figures into comparative and degree constructions, but this is only part of the story. ‘Healthy’ is not just gradable, but also *multidimensional*. There are multiple dimensions or respects of healthiness—musculoskeletal health, freedom from disease, cardiovascular health, and so on—

Contact: Justin D’Ambrosio <jzd1@st-andrews.ac.uk>, Brian Hedden <brian.hedden@anu.edu.au>

and whether Suzy is healthy overall (or all things considered), and whether she is healthier overall than Bill, depend somehow on how healthy Suzy and Bill are in these various respects. More generally, an adjective F is multidimensional if whether its positive form, $\ulcorner$ is $F\urcorner$, applies to an object and whether its comparative forms, $\ulcorner$ is as F as$\urcorner$ and $\ulcorner$ is F-er than$\urcorner$, apply to a pair of objects depend on how those objects stand with respect to multiple underlying dimensions of F-ness.[1]

Multidimensional adjectives are ubiquitous in natural language and crucial to philosophical theorizing [McConnell-Ginet, 1973, Kamp, 1975]. Consider a few more examples: 'democratic,' 'similar,' 'good,' 'rational,' and perhaps even 'conscious.'[2] One country can be more democratic than another with respect to the freedom and fairness of its elections but not with respect to protection of basic liberties. Two possible worlds can be similar with respect to laws of nature but not with respect to the distribution of matters of particular fact. A state of affairs can be good except with respect to inequality, and whether one's beliefs are rational may depend on both how coherent they are as well as how well-proportioned they are to the evidence. Whether and to what extent a creature is conscious may depend on perceptual richness, informational integration, and self-awareness. As these examples show, multidimensionality figures centrally in debates in ethics, epistemology, metaphysics, and beyond, and satisfactory resolution of these debates requires a semantics of multidimensional adjectives.

A key issue concerning the semantics of multidimensional adjectives, and the main issue that will occupy us here, is *the problem of dimensional aggregation*: How do the meanings of the positive and comparative forms of a multidimensional adjective depend on its underlying dimensions? More precisely, how does the location of an object along the underlying dimensions of an adjective F determine whether that object is F overall, or all things considered? And how does the relative location of two objects along these dimensions determine whether one is at least as F as the other overall, or all things considered? To answer these ques-

[1] We confine our discussion to multidimensional adjectives, but arguably other kinds of lexical items—including nouns, verbs, adverbs, and even modals—can be multidimensional as well. See Sassoon [2013b] for an account of multidimensionality in the nominal domain. We leave it to future research to extend our account to other lexical categories.

[2] For the multidimensionality of consciousness, see Bayne et al. [2016], Birch et al. [2020], Lee [forthcoming], D'Ambrosio and Stoljar [ms].

tions, we need to say how the underlying dimensions of *F*-ness can be *aggregated* to yield verdicts about overall or all-things-considered *F*-ness.

In this paper, we develop a semantic framework that explicitly incorporates aggregation functions, where an aggregation function is a function from a set of orderings of objects along the underlying dimensions of *F*-ness to a ranking of objects in terms of overall *F*-ness. This framework allows us to exploit results from social choice theory to address the problem of dimensional aggregation. The problem of dimensional aggregation is analogous to the problem of preference aggregation studied in social choice theory, and social choice theory has yielded a range of technical results showing how many and what kinds of aggregation functions are consistent with various sets of constraints. Our framework allows us to transpose these constraints into the semantic setting and bring these results from social choice theory to bear on the problem of dimensional aggregation.

Unlike standard approaches, our approach allows multidimensional adjectives to be vague not only in their positive form, but also in their comparative form. For us, a multidimensional comparative is vague just in case there are contexts in which there are multiple admissible aggregation functions, where admissibility is determined by axioms governing dimensional aggregation together with further features specific to the context. Moreover, comparative vagueness also generates a novel source of vagueness for the adjective's positive form.[3] While we take no official stand, the results we survey make it plausible that multidimensional comparatives are indeed often vague.

Our framework also allows us to distinguish comparative vagueness from structural features of comparatives such as completeness and transitivity, and to model them independently. In our framework, vagueness has to do with the number of admissible aggregation functions, while completeness and transitivity have to do with features of the orderings such functions output. Our model thus allows for vagueness with or without completeness or transitivity and for completeness or transitivity with or without vagueness, and even for the novel possibility that it is vague whether completeness or transitivity hold. This is important, for philosophers disagree over completeness, transitivity, and

[3]We do not aim to provide a theory of vagueness in general, but only an account of vagueness which has its source in multidimensionality and the availability of multiple candidate ways of aggregating dimensions.

their relation to vagueness. Some argue that multidimensional comparatives like 'at least as good/preferable/rational as' are incomplete, non-transitive, or both, with radical implications for ethics, decision theory, and epistemology [Chang, 2002, Temkin, 2012, Schoenfield, 2012]. Others argue that incompleteness or non-transitivity can be explained away by appeal to vagueness [Broome, 2004, Elson, 2017, Thomas, 2022, Dorr et al., Forthcoming]. Our framework allows us to model and evaluate these competing views, and to provide novel considerations in favour of the latter.

The paper proceeds as follows. In §2 we discuss extant approaches to the semantics of gradable adjectives and the obstacles they face when confronted with multidimensionality. In §3 we present our semantic framework, which explicitly incorporates dimensions and aggregation functions. In §4, we spell out the analogy between preference aggregation and dimensional aggregation and show how to bring the results of social choice theory to bear on the problem of dimensional aggregation. In §5 we compare our proposal with previous work in semantics and philosophy, and in §6 we offer a brief conclusion.

2. Delineation and Degree-Theoretic Approaches

Work on gradable adjectives is standardly divided into two approaches: the delineation approach and the degree-theoretic approach. Here, we briefly survey both and discuss the obstacles they face in dealing with multidimensionality.

2.1. The Delineation Approach

Pioneered by Lewis [1970], McConnell-Ginet [1973], Kamp [1975], and Klein [1980], the delineation approach derives the comparative form of an adjective from its positive form by adopting a form of supervaluationism. For ease of exposition, we focus on Kamp's version of the approach.

Kamp starts with the positive form of a vague adjective, $\ulcorner$is $F\urcorner$. In an ordinary context, $\ulcorner$is $F\urcorner$ is associated with a positive extension (those objects to which it determinately applies), a negative extension (those to which it determinately does not apply), and an extension gap (those to which neither it neither determinately applies nor determinately doesn't apply). Various *admissible precisifications* of the vague adjective then reduce this extension gap by assigning objects therein to either the positive or negative extension of the adjective.

The delineation approach derives the meaning of the comparative from the positive form by quantifying over these precisifications. The basic proposal is that x is at least as F as y just in case the set of admissible precisifications on which x is in the positive extension of F is a superset (possibly improper) of the set of admissible precisifications on which y is in the positive extension of F. More simply, one object is at least as F as another just in case there are no admissible precisifications on which the latter counts as F but the former doesn't.[4]

When applied to multidimensional adjectives, the delineation approach confronts two related problems. First, it entails that while unidimensional comparatives satisfy completeness, multidimensional comparatives do not. Second, it entails that while multidimensional adjectives may be vague in their positive form, they are sharp in their comparative form. We take these up in turn.

An adjective F satisfies completeness just in case for all x and y, either x is at least as F as y or y is at least as F as x.[5] On the delineation approach, unidimensional comparatives satisfy completeness, but multidimensional ones do not. Suppose that x has poor musculoskeletal health but great cardiovascular health, while for y it is the reverse. In many contexts it will be admissible to precisify 'healthy' in such a way that x counts as healthy but y doesn't, or in such a way that y counts as healthy but x doesn't. On the former precisification, cardiovascular health is weighted more heavily than musculoskeletal health, while on the latter the weighting is reversed. The set of admissible precisifications where x counts as healthy will then be neither a subset nor a superset of the set of those where y does. Then, on Kamp's theory, it will be false that x is at least as healthy

[4]There is a problem lurking. If x and y are already both in the positive extension or both in the negative extension of F, the account says that neither is F-er than the other since there are no precisifications of F on which only one of x and y is in the positive extension or the negative extension of F. Kamp addresses this problem by considering 'inadmissible' precisifications fail to include some determinately F objects in the positive extension (or fail to include some determinately non-F objects in the negative extension). Klein [1980] aims to solve this problem by instead appealing to comparison classes: relative to a comparison class that consists of all of those objects that were originally in the positive extension, not all of them will count as F. Then, the claim is that one object is at least as F as another just in case the set of comparison classes relative to which the one counts as F is a superset (possibly improper) of the set of comparison classes relative to which the latter counts as F.

[5]Following Dorr et al. [Forthcoming], we should clarify that this disjunction is only supposed to hold for objects which are at least as F as themselves, so that it wouldn't be a category mistake to talk about their relative F-ness.

as y and also false that y is at least as healthy as x. Thus, on the delineation view, multiple admissible ways of aggregating dimensions typically result in incompleteness.[6]

But we think that this poses a problem for the delineation view. As noted above, some philosophers reject completeness, with radical implications for ethics, decision theory, and epistemology. But they are in the minority. Many philosophers, ourselves included, think that apparent incompleteness, stemming from multiple ways of aggregating dimensions, is really just vagueness [Broome, 2004, Elson, 2017, Dorr et al., Forthcoming] That is, when there are multiple admissible ways of aggregating dimensions, it is determinately the case that either x is at least as F as y or y is at least as F as x, though it may be vague or indeterminate which it is. (Compare how some theorists respond to the Sorites not by rejecting excluded middle, but rather by saying that while it is determinate that everything is either F or $\neg F$, it can be indeterminate which.)

This brings us to the second problem for the delineation approach: it does not allow for multidimensional comparatives to be vague. Rather, the delineation approach treats all apparent vagueness in the comparative as really yielding incompleteness. Whenever there are multiple admissible ways of weighting and aggregating dimensions of F-ness that yield precisifications that disagree about whether two objects are F overall, the delineation approach says that ⌜at least as F as y⌝ is incomplete. There is no room for that comparative to be vague.

Now, some have denied that comparatives can be vague, holding instead that vagueness can only arise in the positive form.[7] But we think that multidimensional comparatives are vague, regardless of whether they satisfy completeness.[8]

[6] Klein [1980] rejects completeness for similar reasons. On his account, two objects are *non-distinct* with respect to F iff there is no comparison class relative to which one counts as F but the other doesn't, and they are *equivalent* with respect to F at c iff there is no context c' more determinate than c such that they are distinct with respect to c'. Equivalence corresponds to the 'exactly as F as' relation. But for multidimensional adjectives, Klein allows for non-distinctness without equivalence, which corresponds to incompleteness.

[7] See Cooper [1995] for this claim, and see Keefe [2000, 12-4] for a persuasive rebuttal.

[8] Indeed, we think it is plausible that an adjective is multidimensional if and only if its comparative form is vague. While this biconditional is plausible, there are potential counterexamples to both of its directions. One possible counterexample to the left-to-right direction is 'logically strong'. Logical strength is multidimensional, with infinitely many underlying dimensions (namely, whether the theory entails the sentence, for each of the infinitely many sentences in the language), but its comparative form 'logically stronger than' is sharp. We are skeptical of this counterexample, since we're skeptical that the comparative

After all, multidimensional comparatives have all of the features that, acccording to Smith [2008] and others, are indicative of vagueness: Sorites susceptibility, blurred boundaries, and borderline cases.

To see this, suppose that Bill has great musculoskeletal health but mediocre cardiovascular health. Then consider a Sorites sequence of possible variants of Suzy: Suzy_1, Suzy_2, $\ldots\text{Suzy}_n$. Each Suzy_i has the same musculoskeletal health, which is somewhat worse than that of Bill. But each has slightly better cardiovascular health than her predecessor, with Suzy_1 having worse cardiovascular health than Bill and Suzy_n having much better cardiovascular health than Bill. We then get a Sorites paradox for 'at least as healthy as': offhand, Suzy_1 is not at least as healthy as Bill, while Suzy_n is at least as healthy as Bill, and yet it seems that for all i, if Suzy_i is not at least as healthy as Bill, then neither is Suzy_{i+1}. Relatedly, there seems to be a blurred boundary between the Suzy_i's that are at least as healthy as Bill and those that aren't—there is no clear cutoff to be found. And as you would expect with such a blurred boundary, many of the Suzy_i's are borderline cases, such that it is indeterminate whether Suzy_i is at least as healthy as Bill.

In sum, the delineation approach errs in treating multiple admissible ways of aggregating dimensions as yielding incompleteness, and in letting incompleteness crowd out vagueness.[9] By contrast, we are sympathetic to completeness and

'at least as logically strong as' is really formed compositionally from a positive form 'logically strong' and the comparative morpheme (compare Dorr et al. [Forthcoming]). Keefe [2000, 12-3] suggests counterexamples to the right-to-left direction; a unidimensional comparative like 'taller than' can be vague, since if it is indeterminate whether Suzy's hair counts as part of her, and she's taller than Bill with the hair included but not without, then it's indeterminate whether Suzy is taller than Bill. We are likewise skeptical of this counterexample, since the indeterminacy is traceable to the vagueness of the name 'Suzy'; once we replace names with sharp noun phrases, any indeterminacy disappears.

[9]Delineation theorists might try to allow for vague comparatives by holding that the metalanguage—and in particular the term 'admissible'—is itself vague (see Williamson [1994, Ch. 5] for discussion). Then it may be that on one (meta-admissible) precisification of 'admissible,' the set of admissible precisifications where x counts as healthy is a proper superset of the set of admissible precisifications where y counts as healthy. But there may be another (meta-admissible) precisification of 'admissible' on which this superset relation does not hold. In that case, it will be vague whether x is healthier than y. But we find this approach objectionable, for it treats the vagueness of multidimensional comparatives quite differently from the vagueness of adjectives in their positive form. The latter is represented in the object language, via the existence of admissible precisifications that differ on which objects count as F. The former is represented only in the metalanguage, via the existence of some precisifications of 'admissible' where certain set-theoretic relations hold between certain sets of admissible precisifications,

to comparative vagueness. But we needn't insist on these views, for the semantic framework we develop in §3 is neutral with respect to them.

2.2. *The Degree-Theoretic Approach*

In contrast to the delineation approach, the degree-theoretic approach [Cresswell, 1973, Kennedy, 1997, Kennedy and McNally, 2005, Kennedy, 2007] takes *scales*, or ordered sets of degrees, as primitive, and uses them to derive the meanings of both the positive and comparative forms of gradable adjectives. On the degree-theoretic approach, the semantic value of a gradable adjective is a function of type $\langle e, d \rangle$ from objects (or *e*ntities) to degrees. So, for instance, the semantic value of 'tall' is a function from objects to their (maximal) degree of height.

In order to derive the meaning of the comparative form from the degree-function denoted by a particular adjective, the degree-theoretic approach makes use of a morpheme called *deg*, which relates the degrees of two objects:

$$[\![deg]\!] = \lambda x \lambda y \lambda F.[F(x) \succsim_F F(y)]$$

This morpheme combines with an adjective F and two objects x and y and returns true just in case $F(x)$ is at least as great as $F(y)$. So, for instance, ⌜x is at least as tall as y⌝ is true just in case x's degree of tallness is at least as great as that of y.

In order to derive the meaning of the positive form, the degree-theoretic approach employs another morpheme called *pos*, which relates an object's degree to some contextually specified standard degree:

$$[\![pos]\!] = \lambda F.\lambda x.[F(x) \succsim_F d^s_F]$$

Using *pos*, the degree theorist holds that ⌜x is tall⌝ is true just in case x has a degree of tallness at least as great as some contextually specified standard d^s_F.

The degree-theoretic approach provides an attractive account of unidimensional adjectives, but what about multidimensional ones? Degree theorists have mostly avoided discussing multidimensional adjectives. Kennedy [2007, 6] briefly mentions them but suggests that they are really just polysemous. He writes that an adjective like 'large,' as applied to cities, 'can be used (at least) to measure either population or sprawl, resulting in different truth conditions.' But both uses

---

and other precisifications of 'admissible' where these set-theoretic relations do not hold.

result in a unidimensional reading of 'large.' We disagree: 'large' can be used in such a way that whether it holds of some city depends on *both* population and sprawl, along with some way of aggregating these dimensions.

Multidimensional adjectives pose *prima facie* problems for the degree-theoretic approach. Standard versions thereof simply assume that we have, for each gradable adjective F, a set of degrees of F-ness and a unique total (i.e. reflexive, transitive, complete, and anti-symmetric) ordering of those degrees. But as noted above, completeness is controversial, and even transitivity has been contested [Rachels, 1998, Temkin, 2012]. While we are sympathetic to both completeness and transitivity, we think it is a virtue of a semantic framework for it to be neutral with respect to these structural features. More problematically, the existence of a uniquely privileged ordering of degrees is incompatible with vagueness in the comparative form, and as we argued in §2.1, multidimensional comparatives seem to exhibit the features of vagueness. Lastly, the degree-theoretic approach simply posits orderings of degrees, whereas we would like such orderings to be derived by somehow aggregating underlying dimensions.

But while multidimensional adjectives present challenges for standard versions of the degree-theoretic approach, we think that they can be modified so as to deal with multidimensionality in a natural and elegant way. Indeed, with some modest additional assumptions discusses at the end of the next section, the semantic framework we advocate can be made degree-theoretic in a way that allows for comparative vagueness, does not presuppose completeness or transitivity, and incorporates aggregation functions. We turn to that framework now.

3. A Semantics for Multidimensional Adjectives

In this section, we develop a semantic framework for multidimensional adjectives that explicitly incorporates dimensions and aggregation functions. In the next section, we note a close analogy between the problem of dimensional aggregation and the problem of preference aggregation in social choice theory, and we show how results from social choice theory can be transposed into our framework to allow us to make progress on the problem of dimensional aggregation.[10]

[10] We are not the first to bring aggregation functions into semantics, nor are we the first to explore the analogy between aggregation of preferences and aggregation of underlying

We start with a function DIM that tells us what the underlying dimensions of a given adjective are and how objects rank along those dimensions. $DIM(F, c, w)$ that takes an input a multidimensional adjective F, context c, and world w and yields as output a *profile of weak orderings* $\langle \succsim_{F_1}, \ldots, \succsim_{F_n} \rangle$ of the domain O of objects relevant in that context. (A weak ordering is reflexive, transitive, and complete.) Each ordering $\succsim_{F_i}$ says how objects rank on underlying dimension i of F.[11]

Since each dimensional ordering $\succsim_{F_i}$ is a weak ordering, it can—given a few further assumptions[12]—be represented by a dimensional value function $V_{F_i} : O \rightarrow \mathbb{R}$ from objects in the domain to real numbers such that $V_{F_i}(x) \geq V_{F_i}(y)$ iff $x \succsim_{F_i} y$. As detailed in §4, there will typically be multiple different dimensional value functions that meet this condition; which value functions—and, more generally, which *profiles* (or vectors) of value functions—represent the same underlying facts, i.e. are 'informationally equivalent', will depend on our assumptions about the measurability and comparability of the different dimensions. Value functions are formally identical to degree functions, at least when degree functions are treated as mapping objects to real numbers. Thus, given a value function V_{F_i} that represents the dimensional ordering $\succsim_{F_i}$, we can treat $V_{F_i}(x)$ as the *degree* to which x is F along dimension i, and treat V_{F_i} as a degree function of type $\langle e, d \rangle$, just as on the standard degree-theoretic approach outlined above.[13]

We next need to consider ways of aggregating dimensional value functions in order to rank objects in terms of overall F-ness. A *dimensional aggregation function* $a : V^n \rightarrow \wp(O^2)$ is a function that takes a profile of n dimensional value functions $\vec{v} = \langle V_{F_1}, \ldots, V_{F_n} \rangle$ as input and returns an 'overall' or 'all-things-considered' ordering $\succsim^a_{\vec{v}}$ of the objects in the domain O.[14] Hence $a(\vec{v}) = a(\langle V_{F_1}, \ldots, V_{F_n} \rangle) = \succsim^a_{\vec{v}}$. Given a multidimensional adjective F and the set of its dimensions relevant in context c, there will be a set $ADM(F, c)$ of such aggregation functions that the context does not determinately rule out; those aggregation functions are *admissible* and the rest are *inadmissible*.

We then need to specify the semantics of both the positive and comparative

dimensions of a multidimensional adjective. In §6 we discuss related work from semanticists and philosophers and how our own approach builds and improves upon this earlier work.

[11] Compare Grinsell [2017, Ch. 3].

[12] See Debreu [1954] for details and discussion.

[13] See Cresswell [1976] and Kennedy [2001] for further discussion of the nature of degrees.

[14] Here V^n is the set of all n-tuples of value functions, or equivalently, profiles of consisting of n value functions, and $\wp(O^2)$ is the set of all sets of ordered pairs from the domain O.

forms of a multidimensional adjective relative to an aggregation function. Since each such function yields as output an ordering of objects in the domain in terms of overall F-ness, this can serve to underwrite the meaning of the comparative. The most natural way is to treat the comparative form ⌜at least as F as⌝ as denoting (relative to an aggregation function and input profile) the ordering outputted by that aggregation function, $\lambda y.\lambda x[x \succcurlyeq^a_{\vec{v}} y]$. Hence, ⌜$x$ is at least as F as y⌝ is true relative to dimensional aggregation function a and input profile $\vec{v}$ if and only if $x \succcurlyeq^a_{\vec{v}} y$. (Following standard practice, we then say that x is F-er than y just in case x is at least as F as y but not *vice versa*, and x and y are equally F just in case x is at least as F as y and *vice versa*.)

What about the positive form? Here our answer takes a page from degree theorists in introducing an object d in the ordering $\succcurlyeq^a_{\vec{v}}$ that sets the standard for qualifying as F. Thus, if F is a multidimensional adjective, then ⌜ is F⌝ denotes (relative to a and $\vec{v}$) the following property: $\lambda x.[x \succcurlyeq^a_{\vec{v}} d]$. This entails that ⌜x is F⌝ is true relative to aggregation function a and input profile $\vec{v}$ if and only if $x \succcurlyeq^a_{\vec{v}} d$.

We can then say that a sentence is determinately true at context c and world w just in case it is true at w relative to all $a \in ADM(F, c)$. It is determinately false just in case it is false relative to all $a \in ADM(F, c)$. Otherwise, it is neither determinately true nor determinately false.[15]

To see this semantics in action, let's consider the example of the multidimensional adjective 'athletic'. Suppose that the dimensions of athleticism relevant in context c are speed (F_1), agility (F_2), and endurance (F_3). $DIM('athletic', c, w)$ then supplies us with a profile of dimensional orderings $\langle \succcurlyeq_{F_1}, \succcurlyeq_{F_2}, \succcurlyeq_{F_3} \rangle$, which can be represented by the profile of value functions $\vec{v} = \langle V_{F_1}, V_{F_2}, V_{F_3} \rangle$. (That profile of dimensional orderings will also be represented by various other profiles of value functions which are informationally equivalent to $\vec{v}$.) An aggregation function a then takes $\vec{v}$ as input and returns $\succcurlyeq^a_{\vec{v}}$, which is an ordering of the objects in the domain in terms of their overall athleticism. Given this ordering, x is at least as athletic as y relative to a and $\vec{v}$ if and only if $x \succcurlyeq^a_{\vec{v}} y$. Context then specifies

[15]While we have been using terminology often associated with the supervaluationist theory of vagueness [Fine, 1975], we are not committed to that approach. We do not, for instance, hold that truth is determinate truth (i.e. truth on all admissible precisifications), which is characteristic of supervaluationism. Instead, our framework is compatible with a wide range of theories of vagueness, including not only supervaluationism but also epistemicism [Williamson, 1994] and theories involving continuum-many degrees of truth [Smith, 2008], among others.

a designated object d as the standard for athleticism, such that ⌜x is athletic⌝ is true (relative to a and $\vec{v}$) if and only if $x \succsim^a_{\vec{v}} d$.

Our framework allows us to model comparative vagueness and structural properties (like transitivity and completeness) independently. Vagueness in the comparative is a matter of there being multiple admissible dimensional aggregation functions in a given context, while transitivity and completeness are matters of the properties of the orderings that the admissible aggregation functions output.[16] A multidimensional comparative is (a) vague and determinately complete (or transitive) iff there are multiple admissible aggregation functions, and each outputs a complete (or transitive) ordering; (b) vague and determinately incomplete (or non-transitive) iff there are multiple admissible aggregation functions, and each outputs an incomplete (or non-transitive) ordering; (c) sharp and complete (or transitive) iff there is one admissible aggregation function, and it outputs a complete (or transitive) ordering; or (d) sharp and incomplete (or non-transitive) just in case there is one admissible aggregation function, and it outputs an incomplete (or non-transitive) ordering. Our framework allows for any of these possibilities. It also allows for the novel possibility that it is vague or indeterminate whether completeness or transitivity hold. This will be the case if some admissible aggregation functions output complete (or transitive) orderings while others output incomplete (or non-transitive) orderings.

This framework provides a more general way of theorising about vagueness than other approaches. First, it preserves the kind of vagueness that philosophers and semanticists typically address in theorising about vague adjectives: vagueness in the standard that an object must meet in order to fall into the extension of the positive form. In adopting the degree-theorist's approach to this standard, our approach allows it to be vague exactly which object d serves as the standard for e.g., height or athleticism, and so allows for borderline cases that result from indeterminacy in this standard. Second, our framework also allows for a further source of vagueness, one that can only arise for multidimensional adjectives: vagueness in how the underlying dimensions are aggregated to yield an ordering of objects in terms of their overall F-ness. The admissibility of mul-

[16] A multidimensional comparative is transitive relative to admissible aggregation a just in case for all input profiles $\vec{v}$ and all $x, y, z \in O$, if $x \succsim^a_{\vec{v}} y$ and $y \succsim^a_{\vec{v}} z$, then $x \succsim^a_{\vec{v}} z$. It is complete relative to a just in case for all input profiles $\vec{v}$ and $x, y \in O$, $x \succsim^a_{\vec{v}} y$ or $y \succsim^a_{\vec{v}} x$.

tiple such orderings yields vagueness in the comparative form, and it also yields a novel source of vagueness in the positive form. Given that an object may meet the contextually specified standard in one admissible ordering but not another, it can be indeterminate whether a given object is at least as F as this standard, and so whether that object is F.

We close this section by considering whether our semantic framework can be considered a degree-theoretic one. Here is one way of making our semantics degree-theoretic. We can take the ordering of objects outputted by our aggregation function and represent it by a value function from objects to real numbers that is unique at least up to strictly increasing transformation. Relative to a choice of one of these informationally equivalent value functions, we can talk about objects' degrees of F-ness (though the ordering of objects would be in some sense more fundamental than the ordering of degrees). But such representation is possible only if the ordering satisfies transitivity and completeness, which we have seen are controversial for multidimensional comparatives.

There is, however, a way in which our semantics can be made degree-theoretic, even without assuming transitivity and completeness. We can think of each object's degree of F-ness as *vector-valued*—i.e. as a vector of its degrees on each underlying dimension, and then use the ordering of objects outputted by an aggregation function to derive an equivalent ordering of their vector-value degrees (though again, the ordering of objects might in some sense be more fundamental than the ordering of degrees).

It doesn't follow from the very definition of an aggregation function that we can do this. It is compatible with that definition that the ordering of objects depends both on how they stand on the underlying dimensions *and* on the nature of those objects. This could happen if e.g., agility plays less of a role than speed and endurance in determining the overall athleticism of horses, while the situation is reversed for basketball players. If that happens, then we may not be able to order objects just by ordering their corresponding vector-valued degrees, for two objects with the same vector-valued degrees could differ in overall athleticism depending on whether they are horses, basketball players, etc. But we can make our semantics degree-theoretic by ruling out this possibility and requiring that if two objects are equally F on each underlying dimension, then they are equally F overall. This constraint comes to us from social choice theory; there, it is known as Pareto Indifference.[17] In the next section, we use social choice theory to explore the semantics of multidimensional adjectives further, exploiting an analogy between aggregating individuals' preferences and aggregating underlying dimensions of a multidimensional adjective.

4. Dimensional Aggregation and Preference Aggregation

Social choice theory is concerned with whether and how it is possible to aggregate individual preferences into an overall 'social' ranking. Call this the *problem of preference aggregation*. There is a deep and illuminating analogy between the problem of preference aggregation in social choice theory and the problem of dimensional aggregation in semantics. This analogy allows us to bring results from social choice theory to bear on the problem of dimensional aggregation.

In social choice theory, we begin with a set of individuals, each with a preference ordering over alternative states of affairs. Given standard assumptions, each preference ordering can be represented by a utility function (unique up to some specified class of transformations), which assigns higher numbers to more preferable states of affairs. The problem of preference aggregation is to determine how the utility functions of individuals can or should be aggregated to yield an ordering of states of affairs in terms of overall, 'social' betterness. Such aggregation is accomplished by a preference aggregation function, which takes as input a profile of utility functions—one per individual—and yields as output an ordering of states of affairs in terms of overall, or all-things-considered, goodness. Social choice theory is then concerned with what constraints should govern preference aggregation and how many and what kinds of functions satisfy them.

In our semantics, we begin with a set of underlying dimensions of an adjective F, each of which provides us with an ordering of a domain of objects with

[17] In social choice theory, an aggregation function is said to satisfy Profile-Dependent Welfarism iff for all input profiles, its output ordering of objects is equivalent to some ordering of their vector-valued degrees. Given transitivity, Profile-Dependent Welfarism is equivalent to Pareto Indifference, which says that if two objects rank equally on every underlying dimension, then they rank equally overall [Blackorby et al., 1990]. In the present context, Profile-Dependent Welfarism would be more aptly called Profile-Dependent *Dimensionalism*, for the underlying dimensions of a multidimensional concept need not be individuals' welfares. See Sen [1977], Blackorby et al. [1990] for further discussion of (different forms of) Welfarism, and see Hedden and Nebel [ms] for discussion and defense of Dimensionalism.

respect to that dimension. Given standard assumptions, each such ordering can be represented by a value function (unique up to some specified class of transformations) which assigns higher numbers to objects that are F-er along that dimension. The problem of dimensional aggregation is to determine how the value functions corresponding to underlying dimensions of F-ness should be aggregated to yield an ordering of objects in terms of overall F-ness. Such aggregation is accomplished by a dimensional aggregation function, which takes as input a profile of value functions—one per underlying dimension—and yields as output an ordering of objects in terms of overall, or all-things-considered, F-ness. Given this background, we are concerned with which constraints govern dimensional aggregation and how many and what kinds of functions satisfy them.[18]

With this analogy in hand, we now turn to surveying some key results from social choice theory and illustrating how they carry over to the problem of dimensional aggregation. We focus on the question of whether the number of admissible aggregation functions for a given adjective in a context is zero, exactly one, or more than one. If 0, then the adjective is incoherent; it is governed by constraints that are jointly unsatisfiable. If 1, then the adjective is sharp in its comparative form; there is some determinately correct aggregation function, and hence some determinately correct ordering of objects in terms of overall F-ness. If more than 1, then the adjective is vague in its comparative form; it is indeterminate how to correctly aggregate its dimensions, and hence indeterminate what is the correct ordering of objects in terms of overall F-ness. We tentatively suggest that the last possibility is the most likely.

[18]While we are impressed by the analogy between preference aggregation and dimensional aggregation, there are also some disanalogies. First, while there may be a single, unified account of rational preference aggregation, it may be that dimensional aggregation works differently for different multidimensional adjectives. Second, some constraints that are plausible in the case of preference aggregation may not be in the context of dimensional aggregation. For instance, each individual's preferences should plausibly count the same in determining overall betterness, but some dimensions of a multidimensional adjective may be weightier than others, as we discuss in §4.2. Third, the problem of preference aggregation is normative. It is concerned with how preferences *ought* to be aggregated. But the problem of dimensional aggregation can be seen as descriptive, normative, or a mixture of the two. We may be concerned with how speakers *in fact* use multidimensional adjectives or with how they *ought* to do so *insofar as they are rational*.

4.1. Zero Admissible Aggregation Functions: Incoherence

The seminal result in social choice theory is Arrow's [1951] impossibility theorem, which shows that a set of plausible constraints on aggregation are jointly unsatisfiable. We begin by stating two 'background' constraints before turning to his 'official' ones. First, Arrow treats individual preferences as merely ordinal and interpersonally non-comparable: while we can talk about whether an individual prefers one thing over another, we cannot talk about *how strong* these preferences are, nor can we talk about whether something would satisfy one person's preferences to a greater degree than another's. In the standard framework of Sen [1970], this assumption is captured with an invariance constraint stating that an aggregation function must output the same overall ordering for any profiles of value functions which are 'informationally equivalent.'[19] Given the assumptions of ordinal measurability and interpersonal non-comparability, two profiles of value functions are informationally equivalent just in case one can be derived from the other by subjecting each value function to some strictly increasing transformation (possibly a different one for each dimension). This yields:

Ordinal Scale Non-Comparability (ONC) :

For all $a \in ADM(F,c)$, $a(\langle V_{F_1}, \ldots, V_{F_n} \rangle) = a(\langle \phi_1(V_{F_1}), \ldots, \phi_n(V_{F_n}) \rangle)$ for any vector of (possibly different) strictly increasing transformations $\langle \phi_1, ..., \phi_n \rangle$

Arrow's second background assumption is that aggregation functions must output weak orderings, which are reflexible, transitive, and complete:

Weak Ordering Outputs (WO) :

For any input profile $\vec{v} = \langle V_{F_1}, \ldots, V_{F_n} \rangle$ in the domain of a,

$a(\langle V_{F_1}, \ldots, V_{F_n} \rangle) = \succsim^a_{\vec{v}}$ is a weak ordering.

In addition to these two background constraints, Arrow's theorem involves four more explicit constraints:

Unrestricted Domain (U) :

a is defined for all logically possible profiles of value functions

[19] This informational invariance framework has recently come under attack from Nebel [2021a,b]. We are sympathetic to his critique, to which we briefly return in §6.

Weak Pareto (P) :

If for all i, $V_{F_i}(x) > V_{F_i}(y)$, then $x \succ^{a}_{\vec{v}} y$

Independence of Irrelevant Alternatives (I) :

If for all i, $V_{F_i}(x) = V^*_{F_i}(x)$ and $V_{F_i}(y) = V^*_{F_i}(y)$, then $x \succcurlyeq^{a}_{\vec{v}} y$ iff $x \succcurlyeq^{a}_{\vec{v}^*} y$

Non-Dictatorship (D) :

There is no i such that whenever $V_{F_i}(x) > V_{F_i}(y)$, $x \succ^{a}_{\vec{v}} y$

Unrestricted Domain requires that aggregation functions be defined for all possible values that dimensional value functions can assign to objects in the domain. Weak Pareto requires that if x ranks strictly above y on all underlying dimensions, then the aggregation function ranks x strictly above y overall. Independence of Irrelevant Alternatives says that the overall ranking of x vis-à-vis y depends only on the values assigned to x and y on each underlying dimension. And Non-Dictatorship says that there is no underlying dimension such that whenever one object ranks higher than another on that dimension, it ranks higher overall.

Given these constraints, we can state Arrow's theorem, adapted to the case of dimensional aggregation: If there is some finite number of underlying dimensions and a finite number ≥ 3 of objects in the domain, then there is no aggregation function that satisfies ONC, WO, U, P, I and D.

Hence if these constraints govern some multidimensional adjective F, then there is no admissible way of aggregating its underlying dimensions to yield an ordering of objects in terms of overall F-ness. This amounts to a kind of *incoherence* in the adjective, whereby we can only talk about whether one thing is F (or whether one thing is F-er than another) *in some respect*, but not about whether one thing is F (or whether one thing is F-er than another) *overall*.

We find this incoherence implausible, at least for most multidimensional adjectives. For we seem to use them felicitously, even without restricting attention to a single dimension. To avoid such incoherence, we must deny that all of the Arrovian constraints determinately govern most multidimensional adjectives.

Which constraint(s) should we reject? We think that Weak Pareto (P) is almost incontestable.[20] Non-Dictatorship (D) also seems reasonable, at least for most

[20] See Hedden and Muñoz [forthcoming] for a defense of Weak and Strong Pareto for value pluralism.

adjectives, because violating it would entail that one underlying dimension is lexically prior to all others. We also find Independence of Irrelevant Alternatives (I) rather plausible, though perhaps a bit less so than (P) and (D).

Unrestricted Domain (U) might be contested on the grounds that some combinations of dimensional values assigned to objects may be logically possible and yet wildly unrealistic, such that we can effectively rule them out in many contexts. With 'intelligent,' for instance, it may be quite unlikely for someone to have extremely high verbal intelligence but extremely poor spatial reasoning, and so we may not be troubled if our aggregation function is undefined for such unrealistic inputs.[21] But even if we reject (U), social choice theory has shown that various weaker domain conditions still yield impossibilities in the presence of the other constraints above [Gaertner, 2001].

We think that the most attractive escape routes from Arrow's impossibility theorem, for the case of dimensional aggregation, are to reject ONC or WO. If we reject WO, we can hold that multidimensional comparatives are non-transitive, incomplete, or both (§4.2). We explore these possibilities §4.2. If we reject ONC, we can hold that underlying dimensions are measurable not just ordinally, but cardinally, and/or we can hold that they admit of certain meaningful interdimensional comparisons of value (§4.3).

4.2. Exactly One Admissible Aggregation Function: Sharpness

Suppose that we are not happy to concede that multidimensional adjectives are incoherent, as the applicability of Arrow's theorem would imply. How might we escape this pessimistic conclusion? Here, we consider the option of abandoning the assumption that admissible aggregation functions must output weak orderings (WO). There are two obvious ways of doing so. The first is to allow overall orderings to be non-transitive. The second is to allow them to be incomplete. We take up these possibilities in turn. Denying each assumption opens up a range of possibilities for aggregation, but in each case, by imposing a collection of further constraints, we can uniquely characterize a natural WO-denying aggregation rule. Then, if these further constraints determinately govern some multidimensional adjective, then there is exactly one admissible aggregation function for

[21]See Hedden and Nebel [ms] for further discussion in the context of multidimensional concepts.

that adjective, with the result that it is sharp in its comparative form.

First, we might reject WO by rejecting transitivity. This opens up the possibility of embracing e.g., Majority Rule:

Majority Rule :

> x is at least as F as y overall if and only if the number of dimensions on which x is at least F as y is at least as great as the number of dimensions on which y is at least as F as x.
> $\rightsquigarrow x \succcurlyeq_F y$ iff $|\{i|V_{F_i}(x) \geq V_{F_i}(y)\}| \geq |\{i|V_{F_i}(y) \geq V_{F_i}(x)\}|$

It is clear that majority rule satisfies U, P, I, and D. But without WO, there are many other aggregation rules that likewise satisfy these conditions. Accordingly, if we give up only WO while retaining the rest of Arrow's constraints, there will be multiple admissible ways of aggregating their dimensions in context—instead of being incoherent, multidimensional adjectives will be vague.

However, May [1952] proved that majority rule is the only aggregation function that satisfies the following four constraints together with ONC: Unrestricted Domain, Neutrality, Anonymity, and Positive Responsiveness.[22] Unrestricted Domain is as before. Neutrality says that there is nothing special about any of the objects in terms of how they are treated by the aggregation function. In particular, if two pairs of objects $\{x,y\}$ and $\{w,z\}$ are alike in terms of how the first member of the pair is ranked vis-à-vis the second along each underlying dimension, then they must be alike in terms of how the first is ranked vis-à-vis the second overall. Anonymity says that there is nothing special about any of the underlying dimensions in terms of how they are treated by the aggregation function; if we permute the value functions of the underlying dimensions, so that two dimensions 'swap' their value functions, then the overall ordering remains the same. Positive Responsiveness says that if x and y are originally equally F overall, and if one underlying dimension's ranking shifts in favor of x vis-à-vis y (i.e. ranking x above y where they were previously ranked equally, or ranking them equally where previously y was above x), then x is now F-er than y overall.

[22] May was working in a framework where the inputs to aggregations functions are profiles of weak orderings, rather profiles of value functions. But May's theorem will also apply in the latter framework once we add the assumption of ordinal non-comparability.

It is well-known that Majority Rule violates the requirement that $\succcurlyeq$ be transitive. Indeed, it even violates the weaker requirement that $\succ$ be acyclic.[23] After all, it can happen that a majority of dimensions rank x above y, a (different) majority of dimensions ranks y above z, and a (different) majority of dimensions ranks z above x. In this case, majority rule says that x is F-er than y, y is F-er than z, and z is F-er than x. This is Condorcet's paradox. Some theorists, most notably Rachels [1998] and Temkin [2012], reject transitivity and even acyclicity for the relation 'better than.' If they are correct, then transitivity and acyclicity might be rejected for other multidimensional comparatives as well. But still, Rachels and Temkin hold a minority view. Most think that all comparatives necessarily obey transitivity and acyclicity [Broome, 2004, 2013, Nebel, 2018].

Second, might reject WO by rejecting completeness rather than transitivity. This opens up the possibility of embracing e.g., the Strong Pareto Rule:

Strong Pareto Rule :

> x is at least as F as y overall if and only if x is at least as F as y on each underlying dimension.
> $\rightsquigarrow x \succcurlyeq_F y$ if and only if for all i, $V_{F_i}(x) \geq V_{F_i}(y)$

Weymark [1984, Theorem 3] gives a unique characterization of the Strong Pareto Rule. He proves that it is the only aggregation function that outputs quasi-orderings (i.e. relations which are reflexive and transitive, but not necessarily complete) and satisfies ONC[24], U, I, A, and the Strong Pareto *Principle* P*, which says that if x ranks at least as high as y on all dimensions, then x ranks at least as high as y overall, and if, moreover, x ranks strictly higher than y on some dimensions, then x ranks strictly higher than y overall.[25]

The Strong Pareto Rule leads to rampant incompleteness. For it says that if x is F-er than y on some underlying dimensions but less F on others, then neither is

[23]Acyclicity is the requirement on orderings F that if $x \succ_F y$ and $y \succ_F z$, then $x \succcurlyeq_F z$.

[24]Like May, Weymark was working in a framework where the inputs to aggregation functions are profiles of weak orderings. But his theorem will still apply in the framework where the inputs are profiles of value functions, provided we add the assumption ONC.

[25]See also Sen's [1970, Theorem 5.3] unique characterization of the Pareto Extension Rule, which is like the Strong Pareto Rule except that it replaces any incompleteness with indifference: x is at least as F as y just in case x is F-er than y on at least one underlying dimension. Sen shows that this is the only aggregation function that satisfies U, P*, I, and A and which outputs reflexive and quasi-transitive orderings. (Quasi-transitivity is transitivity of $\succ_F$).

at least as F as the other overall. Now, we have already seen that there is a lively debate about whether multidimensional comparatives obey completeness, and we are officially neutral on completeness for the purposes of this paper. Having said that, even if multidimensional comparatives need not satisfy completeness across the board, we suspect that they do not typically yield the kind of pervasive incompleteness that would result from the Strong Pareto Rule.

Still, whether we reject transitivity (making room for Majority Rule) or completeness (making room for the Strong Pareto Rule), neither of the above uniqueness results seems particularly relevant to the semantics of multidimensional adjectives. This is because Anonymity is implausible as a constraint on dimensional aggregation. In the context of social choice, Anonymity is motivated by fairness; it shouldn't matter *who* has which preferences. But in the context of multidimensional adjectives, no such ideal of fairness to dimensions applies; some might be weightier than others. We are not ruling out the possibility that some constraints might fix a unique rule for dimensional aggregation, but the most well-known uniqueness results involve constraints like Anonymity that are implausible in the case of dimensional aggregation.

4.3. *More than One Admissible Aggregation Function: Vagueness*

So far, we have examined two possibilities: one in which multidimensional adjectives are governed by jointly unsatisfiable constraints and are therefore incoherent, and another in which they are governed by uniquely satisfiable constraints and are therefore sharp in their comparative forms. Here, we look at the third possibility, namely that they are governed by constraints satisfiable by multiple aggregation functions, meaning that they are vague in their comparative forms.

We suggest that this possibility will likely obtain if we reject ONC and allow aggregation functions to take into account more information than just the ordinal rankings of objects along underlying dimensions. Put simply, enriching the informational structure of underlying dimensions avoids incoherence by incurring vagueness. (Rejecting ONC is not the only route to vagueness, however, as we noted in the previous subsection and reiterate at the end of this one.)

First, suppose that our underlying dimensions are measurable on interval scales, which have strictly richer structure than ordinal scales. Here, both the order of the objects and the ratios of gaps between numbers assigned to different objects are meaningful. But the choice of a unit and origin (zero point) is not. This means that if a value function V represents a dimension's ranking, then so do all and only positive affine transformations $V^* = aV + b$ $(a > 0)$ thereof. A standard example of an interval scale-measurable dimension is temperature. It can be represented by both Fahrenheit and Celsius scales, which differ in their choice of both unit and origin but agree on ratios of temperature differences.

Suppose also that certain 'interdimensional comparisons of value,' analogous to interpersonal comparisons of utility, are possible. In particular, suppose that our underlying dimensions are 'unit comparable,' so that for any x, y, z, u, we can compare the degree to which x ranks above y on one dimension with the degree to which z ranks above u on another dimension. So it is meaningful to say, e.g., that the degree to which x is healthier than y in terms of cardiovascular fitness is equal to the degree to which z is healthier than u in terms of cholesterol. On Sen's [1970] approach, these two assumptions—interval scale measurability and interdimensional unit comparisons—yield the following invariance condition:

Interval Scale Unit Comparability (IUC) :

For any vector of positive affine transformations $\langle \phi_1, ..., \phi_n \rangle$ with common unit (i.e. where each ϕ is such that $\phi_i(V) = aV + b_i$, with $a > 0$)

$$a(\langle V_{F_1}, ..., V_{F_n} \rangle) = a(\langle \phi_1(V_{F_1}), ..., \phi_n(V_{F_n}) \rangle)$$

Sen [1970, Theorem 7*1] shows that there are aggregation functions which satisfy IUC along with WO, U, P, I, and D. One particularly natural example is a 'utilitarian' weighted sum aggregation function, where all weights c_{F_i} are non-negative:

Utilitarian Aggregation :

$$x \succcurlyeq y \text{ iff } c_{F_1} V_{F_1}(x) + ... + c_{F_n} V_{F_n}(x) \geq c_{F_1} V_{F_1}(y) + ... + c_{F_n} V_{F_n}(y)$$

If all and only IUC, WO, U, P, I, and D determinately govern the semantics of some multidimensional adjective, then we get vagueness in the comparative. For one thing, we have not uniquely characterized utilitarian aggregation; it is not the only type of aggregation function that satisfies these constraints. For another, utilitarian aggregation is a functional *form*, not a particular aggregation function. We have specified neither the weights c_{F_i} nor the particular interdimensional unit comparisons.[26] Unless we can uniquely specify these weights and interdimensional comparisons, we are left with multiple—indeed, uncountably many—admissible aggregation functions all of which are utilitarian in form.

Second, suppose our underlying dimensions are measurable on ratio scales. Ratio scales have strictly richer structure than interval scales. Here, the location of the origin is also meaningful, in addition to the order of objects and the ratios of gaps between numbers assigned to them. But the choice of a unit is arbitrary. This means that if V represents a ranking of objects along some dimension, then so do all and only similarity transformations thereof; i.e. all and only those V^* such that $V^* = a \times V$ $(a > 0)$. Examples of ratio scale-measurable dimensions include length, mass, and volume. This means that a multidimensional adjective like 'large,' with underlying dimensions of e.g., length, mass, and volume, may have underlying dimensions which are all ratio-scale measurable.

Now, the very structure of ratio scales means that some interdimensional comparisons are meaningful. In particular, all comparisons of percentage differences along different dimensions are meaningful, as are comparisons of levels across dimensions where the objects are assigned values with different signs (positive, negative, or zero) by the different dimensions. For these percentage differences and signs are preserved by similarity transformations. But with no *other* interdimensional comparability, we get the following invariance condition:

Ratio Scale Non-Comparability (RNC) :

For any vector of similarity transformations $\langle \phi_1, ..., \phi_n \rangle$

$$a(\langle V_{F_1}, ..., V_{F_n} \rangle) = a(\langle \phi_1(V_{F_1}), ..., \phi_n(V_{F_n}) \rangle)$$

With RNC, there are aggregation functions which satisfy intuitively compelling constraints, provided that all dimensional value functions assign only non-negative values. This non-negativity condition seems to be met for 'large,' since it is impossible to have negative values along the underlying dimensions of e.g., length, mass, and volume. Given this non-negativity condition, Tsui and Weymark [1997, Theorem 5] show that an aggregation function satisfies RNC,

[26] d'Aspremont and Gevers [1977] give a unique characterization of this 'utilitarian' functional form, but their conditions include Anonymity, whose implausibility for multidimensional adjectives we have already noted. And they do not give any unique specification of the interdimensional unit comparisons.

WO, WP, and a technical Continuity condition[27] if and only if it has the functional form of Cobb-Douglas aggregation with non-negative exponents:

Cobb-Douglas: $x \succcurlyeq y$ iff $V_{F_1}(x)^{c_{F_1}} \times \ldots \times V_{F_n}(x)^{c_{F_n}} \geq V_{F_1}(y)^{c_{F_1}} \times \ldots \times V_{F_n}(y)^{c_{F_n}}$

If all and only Tsui and Weymark's constraints govern the semantics of some multidimensional comparative, then we again have vagueness.[28] For we have no uniquely privileged assignment of weights to dimensions, meaning that we have multiple (indeed, uncountably many) admissible aggregation functions which share the function form of Cobb-Douglas aggregation but differ in their exponents c_{F_i}.

We think it highly plausible that multidimensional comparatives are indeed vague. We have independent grounds for thinking this, namely that multidimensional comparatives seem to display the features typically regarded as characteristic of vagueness: Sorites susceptibility, blurred boundaries, and borderline cases (§2.1). But we also think that ONC should be rejected for many multidimensional adjectives, and we have seen that rejecting ONC leads naturally to the view that there are multiple admissible aggregation functions. Rejecting ONC is not the only way to generate comparative vagueness, however. First, every other Arrovian constraint is such that if we reject it, there will be multiple aggregation functions compatible with the remaining constraints, so we will still get comparative vagueness unless we also impose further non-Arrovian constraints, as we explored in §4.2. Second, it may be indeterminate which of Arrow's constraints fails for a given multidimensional comparative, again leaving us with multiple aggregation functions that are not determinately ruled out.

5. Related Proposals

We are not the first to investigate aggregation functions in the context of multidimensional adjectives or to view multidimensionality through the lens of social

[27]Continuity says that for all vectors u of dimensional value functions in the domain of such vectors D, the sets $\{v \in D \mid v \succcurlyeq u\}$ and $\{v \in D \mid u \succcurlyeq v\}$ are closed. Intuitively, continuity requires that if one vector of values is better than another, then another vector resulting from a 'sufficiently small' change to the one is still better than the other.

[28]However, they also give an impossibility theorem: there is no aggregation function that satisfies these constraints along with Non-Dictatorship (D) if some dimensional value functions assign negative values.

choice theory. In this section, we discuss how our work builds and improves upon previous work in semantics and philosophy.

A number of semanticists have discussed multidimensional adjectives and aggregation, but without drawing upon work in social choice theory. Sassoon [2013a,b, 2015], working within a broadly degree-theoretic framework, proposes that each multidimensional adjective F is associated with a contextually specified set of underlying dimensions, each of which can be represented by a degree function that maps each object to its maximal degree along that dimension. So far, her framework is rather similar to ours.

Sassoon then confronts the problem of dimensional aggregation. Starting with the positive form, she says that there is some contextually specified threshold for each dimension, and that aggregation works by quantifying over the dimensions on which a given object meets that threshold. But different adjectives employ different kinds of quantification; some employ universal quantification, others employ existential quantification, and still others employ 'dimension counting' [Sassoon, 2013a,b, 2015]. 'Healthy' is of the first type—someone is healthy iff they meet the contextually specified threshold on *all* underlying dimensions. 'Sick' is of the second type—someone is sick iff they meet the contextually specified threshold on *some* underlying dimension(s). 'Intelligent' is of the third type—someone is intelligent iff they meet the contextually specified threshold on some contextually specified number of dimensions.

What about the comparative form? Sassoon's view of the comparative is somewhat difficult to pin down. But she tentatively suggests that the comparative form works via the same type of rule that governs the positive form [Sassoon, 2013a, 368]. Thus, if 'healthy' works via universal quantification, so does 'at least as healthy as': x is at least as healthy as y iff x ranks at least as highly as y on *all* underlying dimensions of health. (This is the Strong Pareto Rule from §4.2, though Sassoon does not make this connection.) If 'sick' works via existential quantification, so does 'at least as sick as': x is at least as sick as y iff x ranks at least as highly as y on *some* underlying dimensions of health. If 'intelligent' works via dimension counting, so does 'at least as intelligent as': x is at least as intelligent as y iff x meets the contextually specified threshold on *at least as many* underlying dimensions as does y [Sassoon, 2015, 18].

We are skeptical of Sassoon's proposals about aggregation; in some cases they yield implausible results. First, any comparative whose dimensions are aggregated via universal quantification (i.e. via the Strong Pareto Rule) will yield rampant incompleteness, since neither x nor y will be at least as healthy as the other whenever one ranks higher on some dimensions but lower on others. While we are officially neutral on whether completeness holds across the board, we have already noted (§4.2) that we find this rampant incompleteness implausible.

Second, any comparative whose dimensions are aggregated via existential quantification will violate the highly compelling Strong Pareto principle P^*, which we saw in §4.2. Aggregation via existential quantification says that x is at least as F as y just in case x ranks at least as highly as y on some dimension. But now let x and y be ranked equally on one dimension and y above x on all others. By the existential aggregation function just mentioned, x will count as at least as F as y, since it ranks equally with y (and hence at least as highly as y) on one dimension. But this conflicts with P^*, which entails that y is F-er than x overall.

Third, any comparative whose dimensions are aggregated via 'dimension counting' will violate the even more compelling Weak Pareto principle P we saw in §4.1. For x and y could meet the contextually specified threshold on exactly the same (and hence the same number of) underlying dimensions as each other (so that the dimension counting rule says that each is at least as F as the other), even though x ranks strictly higher than y on all of them.

A number of other semanticists have discussed multidimensionality and aggregation in connection with subjectivity or judge-dependence. Subjective adjectives are adjectives such as 'tasty' that have an experiential or evaluative component and as a consequence give rise to so-called 'faultless disagreement,' in which two agents disagree about whether the adjective applies without either of them speaking falsely. Some semanticists [Bylinina, 2014, McNally and Stojanovic, 2017, Solt, 2018] have suggested that multidimensional adjectives may be similarly subjective, and even that multidimensionality may be the source of subjectivity, since speakers can disagree about how to aggregate underlying dimensions, and hence disagree about whether a multidimensional comparative holds of two objects, without either speaking falsely.

We agree that there are many contexts in which multiple aggregation functions are admissible, and that in such contexts interlocutors may employ different such functions. But we think that such multiplicity amounts to vagueness in the comparative, and vagueness in general can give rise to apparent subjectivity, as different speakers can precisify an adjective in different ways without making an obvious mistake. But this is, of course, not unique to multidimensionality, for vagueness can arise from other sources as well.

While these other semanticists discuss aggregation but not social choice theory, Grinsell [2017] puts social choice theory at the center of his treatment of multidimensional adjectives. Grinsell first focuses on providing an explanation of the vagueness of multidimensional adjectives in their positive form and then tries to generalize his proposal to vagueness in unidimensional adjectives. He does not address the vagueness of comparatives. In this respect, his aims are quite different from our own, but we nonetheless think it is worth highlighting the points of similarity and difference between our views.

Grinsell's aim is to show that the characteristic features of vagueness are consequences of impossibility theorems like Arrow's. Following Cobreros et al. [2012, 349], Grinsell holds that each vague gradable adjective is associated with an indifference relation and that (as they put it) 'the non-transitivity of the indifference relation is a central feature of all vague predicates.' For Grinsell, as for Cobreros et al., this indifference relation is something like the relation of being *indistinguishable with respect to F-ness*. (Cobreros et al. mention 'not looking to have distinct heights' as the indifference relation for 'tall.') Grinsell's idea is that the inductive premise of the Sorites is motivated by the apparent truth of a principle of tolerance, which says that if any two things are indifferent with respect to F-ness, then either both are F or neither is. According to Grinsell, a theory of vagueness must explain why tolerance is in fact false, despite its appeal.

Grinsell aims to develop a theory that does just this. First, Grinsell treats the semantics of multidimensional adjectives as dependent on an aggregation function that aggregates the orderings associated with each underlying dimension. He also maintains that the Arrovian constraints U, P, I, and D typically govern such aggregation functions (in addition to ONC, which he leaves implicit). Second, Grinsell then argues that any multidimensional adjective F whose aggregation function is governed by these constraints will have a non-transitive indifference relation, $\sim_F$, because in the original Arrovian setting, allowing the indifference relation for preferences to be non-transitive is one (albeit not the only) way to avoid Arrovian incoherence. But the principle of tolerance can be true only if the indifference relation it employs is transitive, since otherwise we could connect something which is clearly F to something which is clearly $\neg F$ through a Sorites sequence of objects, each of which is indifferent to its predecessor. Thus, since Arrow's theorem suggests that the indifference relation may be non-transitive, it helps explain what is arguably a key feature of vagueness.

We welcome Grinsell's introduction of social choice theory into the semantics of multidimensional adjectives, as well as his work in connecting it to vagueness. But we are skeptical of his proposal. The reason is that the indifference relation in social choice theory—i.e. the relation that Arrow's theorem suggests may be non-transitive—is not the same as the indifference relation at work in the principle of tolerance. The indifference relation $\sim$ in social choice theory is defined so as to be incompatible with the preference relation $\succ$ [Sen, 1970]. In particular, we start with the relation $\succcurlyeq$ and define $\succ$ as its asymmetric part ($x \succ y =_{df} x \succcurlyeq y \wedge \neg(y \succcurlyeq x)$) and $\sim$ as its symmetric part ($x \sim y =_{df} x \succcurlyeq y \wedge y \succcurlyeq x$).

In the context of multidimensional adjectives, it is natural to understand $\succ_F$ as corresponding to the relation *F-er than*. This is how we have understood it in our semantics, and as we read Grinsell, it is how he understands it as well [see Grinsell, 2017, p. 73]. But this means we cannot understand $\sim_F$ as corresponding to the relation *indistinguishable with respect to F-ness*, for this relation is not incompatible with the relation of being F-er than. Two things can be indistinguishable with respect to F-ness even though one is in fact F-er than the other, albeit by an amount undetectable through normal means. Instead, the indifference relation $\sim_F$ should be understood as something like 'exactly as F as,' so that it is genuinely incompatible with $\succ_F$, as required. (This interpretation is bolstered by the fact that $\sim_F$ is defined as the symmetric part of $\succcurlyeq_F$; if x and y are each at least as F as the other, then it seems that they must be equally F.) And the 'exactly as F as' relation is necessarily transitive [Broome, 2004, Hedden, 2020].

Having said that, we agree for independent reasons that the relation *indistinguishable with respect to F-ness* is non-transitive (and that the principle of tolerance formulated in terms of it is false), and we have no objection to calling this relation of indistinguishability an indifference relation, so long as we are careful to distinguish it from the indifference relation $\sim_F$ defined in terms of the $\succcurlyeq_F$ relation, to which the results of social choice theory apply.[29] (Note that Cobreros et al.

[29]Moreover, we also think that indistinguishability may be red herring, since the inductive

[2012] do not suggest that their indifference relation is the same as the $\sim$ relation of social choice theory.) Moreover, we don't need social choice theory to explain why indistinguishability is non-transitive; we know that on the basis of more general epistemic considerations about the limits of our knowledge [Williamson, 1990, Hedden, 2020].[30]

Finally, a number of philosophers have used social choice theory to analyze particular philosophically interesting cases of multidimensionality. Hurley [1985] does so for value pluralism in ethics (see also Hedden and Muñoz [forthcoming]); Okasha [2011] for pluralism about theoretical virtues in the philosophy of science; Morreau [2010] and Kroedel and Huber [2013] for overall similarity in the context of counterfactuals; and MacAskill [2016] for decision-making under moral uncertainty. These proposals have provided inspiration for our own, but our work builds upon theirs in several ways. First, these authors are each concerned with a particular case of multidimensionality, whereas we are aiming for a more general approach which brings the results of social choice theory to

premise of the Sorites is compelling even when formulated in terms of small but still noticeable differences, e.g., that if someone with a height of n cm is tall, then so is someone with a height of $n-1$ cm. Pointing out that indistinguishability is non-transitive fails to address why this inductive premise—which makes no reference to indistinguishability, and indeed allows for distinguishability—is false.

[30]Note that Grinsell's proposal also would not apply to the vagueness of unidimensional adjectives in their positive form, though elsewhere [2017, Ch. 4] he suggests that the use of unidimensional adjectives is subject to competing standards, in such a way that Arrow's theorem could apply to attempts to aggregate verdicts from these standards. Grinsell [2012] offers an earlier account of vagueness (for multidimensional adjectives) that is different from, but perhaps complementary to, his [2017] account. The earlier account focuses on Chichilnisky's [1982] theorem, which says that no aggregation function can satisfy WO, P, D, and Continuity (see fn. 28), which says, roughly, that small changes to the input profile should not result in large changes to the output ordering. Grinsell concludes that aggregation functions applying to multidimensional adjectives must be discontinuous, and he connects this to the fact that small changes must sometimes affect whether the positive form of a vague adjective applies to a given object, which might be regarded as a kind of discontinuity. While suggestive, Grinsell's proposal again faces problems. The main one is that the discontinuity at issue in Chichilnisky's theorem is not the same as the 'discontinuity' displayed by vague adjectives. The latter 'discontinuity' is the fact that small differences between objects sometimes yield 'large' differences in whether they are F. The former is the phenomenon whereby small changes to the input profile sometimes yield large changes in the output ordering—i.e. to which objects are F-er than which. In the Sorites paradox, the latter kind of discontinuity is irrelevant, for it does not involve considerations of different ways that the objects might be ordered by the various underlying dimensions of our multidimensional adjective. The Sorites shows that even with a fixed set of objects and a fixed ordering thereof, there must be two objects which are almost exactly as F as each other, but where one is F and the other not.

bear on multidimensional adjectives generally. Second, we are giving a semantics, whereas these philosophers were not concerned with language *per se*. Third, their analyses tended to focus on whether a version of Arrow's theorem applies to a given case of multidimensionality, whereas we also connect social choice theory to comparative vagueness and the structural features of multidimensional comparatives.

6. Conclusion

Many, and perhaps even most, gradable adjectives F are multidimensional; they involve a set of underlying dimensions of F-ness, which must somehow be aggregated to yield verdicts about overall, or all-things-considered, F-ness. Our semantic framework explicitly incorporates aggregation functions, which output an overall ordering of objects that serves as the meaning of a multidimensional comparative. This framework allows multidimensional comparatives to be vague (i.e. if there are multiple admissible aggregation functions), and treats their vagueness as independent of structural features like transitivity and completeness. However, the framework itself is neutral with respect to both.

This framework is a fruitful one, allowing us to bring results from social choice theory to bear on the semantics of multidimensional adjectives. In this regard, our discussion is preliminary, and we think that it points to a number of further avenues of research, two of which we will briefly discuss in closing.

First, we saw in §4.3 that one prominent way to escape from Arrow's impossibility theorem is to hold that certain *interdimensional comparisons of value* are meaningful. Interdimensional comparisons of value are analogous to *interpersonal comparisons of utility* in social choice theory, and economists and philosophers now largely hold them to be meaningful. But in the context of multidimensional adjectives, they are analogous to what are sometimes called *interadjectival comparisons*, which are widely held to be meaningless (for discussion, see Bale [2008], van Rooij [2011]). It seems nonsensical to say e.g., 'This table is heavier than that sofa is heavy.' In recent work, however, [Nebel, 2021a,b, forthcoming] argues that we can aggregate welfares without making interpersonal comparisons of utility.[31] This suggests that we might similarly be able to aggregate

[31] Very briefly, he argues that we should reject the framework of informational invariance

underlying dimensions without making interdimensional comparisons of value, thereby avoiding commitment to nonsensical interadjectival comparisons.

Second, in social choice theory, it is natural to assume that all individuals' preferences are measurable on the same type of scale, though there is considerable debate about what type of scale that is. We have been making the analogous assumption that all underlying dimensions of a multidimensional adjective are measurable on the same type of scale. But while this assumption of common scale types is extremely plausible in the case of social choice theory—people are people, after all—it is far less plausible in the case of multidimensional adjectives. If a multidimensional adjective has underlying dimensions measurable on different types of scales, how if at all can they be aggregated? There is scarcely any work on this problem in social choice theory, the only exception being a series of underappreciated impossibility theorems [Khmelnitskaya, 1999, Khmelnitskaya and Weymark, 2000]. If such impossibility theorems apply in the context of multidimensional adjectives, then we again face the possibility that some such adjectives are incoherent. See Hedden and Nebel [ms] for discussion of how to potentially escape this pessimistic conclusion. For now, we simply mention the problem as another case where the semantics of multidimensional adjectives stands to be illuminated by the results of social choice theory.

---

conditions that is standard in social choice theory and which we employed above. This is because this framework is unable to distinguish between genuine differences in the degrees to which objects instantiate the properties that correspond to the various underlying dimensions, on the one hand, and mere representational changes, on the other [see also Sen, 1977]. In its place, Nebel proposes a qualitative framework in which, roughly, value functions assign not dimensionless real numbers, but rather dimensioned quantities, to alternatives. For instance, a value function representing the dimension of mass would assign me not the number 70, which is dimensionless, but rather the dimensioned quantity 70kg. Nebel shows that in this qualitative framework, a constraint somewhat analogous to the standard informational invariance constraints, which he calls *automorphism invariance*, is unmotivated. And once we reject it, we open up space for a much wider range of aggregation functions than before. We are sympathetic to this work, and we note that adopting Nebel's qualitative framework would only bolster the case that multidimensional comparatives are vague, given the way in which it opens space for more kinds of aggregation functions than in the standard framework.

Acknowledgements

We are grateful to Kevin Dorst, Alan Hájek, Daniel Muñoz, Daniel Stoljar, and an audience at the Dianoia Institute at the Australian Catholic University for their helpful comments and feedback. Authors are listed in alphabetical order and contributed equally to this paper.

References

Kenneth Arrow. A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4):328–346, 1951.

Alan Clinton Bale. A universal scale of comparison. *Linguistics and Philosophy*, 31(1):1–55, 2008.

Tim Bayne, Jakob Hohwy, and Adrian M Owen. Are there levels of consciousness? *Trends in cognitive sciences*, 20(6):405–413, 2016.

Jonathan Birch, Alexandra K Schnell, and Nicola S Clayton. Dimensions of animal consciousness. *Trends in cognitive sciences*, 24(10):789–801, 2020.

Charles Blackorby, David Donaldson, and John A Weymark. A welfarist proof of arrow's theorem. *Recherches Économiques de Louvain/Louvain Economic Review*, 56(3-4):259–286, 1990.

John Broome. *Weighing Lives*. Oxford University Press, New York, 2004.

John Broome. *Rationality Through Reasoning*. Blackwell, London, 2013.

Lisa Bylinina. The grammar of standards. *LOT Dissertation Series*, 2014.

Ruth Chang. The possibility of parity. *Ethics*, 112(4):659–688, 2002.

Graciela Chichilnisky. Social aggregation rules and continuity. *The Quarterly Journal of Economics*, 97(2):337–352, 1982.

Pablo Cobreros, Paul Egré, David Ripley, and Robert van Rooij. Tolerant, classical, strict. *Journal of Philosophical Logic*, 41(2):347–385, 2012. doi: 10.1007/s10992-010-9165-z.

Neil Cooper. Paradox lost: Understanding vague predicates. *International Journal of Philosophical Studies*, 3(2):244–269, 1995.

M. J. Cresswell. The semantics of degree. In Barbara Partee, editor, *Montague Grammar*. Academic Press, New York, 1976.

Max Cresswell. *Logics and Languages*. Methuen and Co. Limited, London, 1973.

Justin D'Ambrosio and Daniel Stoljar. The indeterminacy of consciousness. ms.

Claude d'Aspremont and Louis Gevers. Equity and the informational basis of collective choice. *The Review of Economic Studies*, 44(2):199–209, 1977.

Gérard Debreu. Representation of a preference ordering by a numerical function. *Decision processes*, 3:159–165, 1954.

Cian Dorr, Jacob Nebel, and Jake Zuehl. The case for comparability. *Noûs*, Forthcoming.

Luke Elson. Incommensurability as vagueness: A burden-shifting argument. *Theoria*, 83(4):341–363, 2017. doi: 10.1111/theo.12129.

Kit Fine. Vagueness, truth, and logic. *Synthese*, 30(3/4):265–300, 1975.

Wulf Gaertner. *Domain Conditions in Social Choice Theory*. Cambridge University Press, Cambridge, 2001.

Timothy Grinsell. Avoiding predicate whiplash: social choice theory and linguistic vagueness. *Proceedings of SALT*, 22:424–440, 2012.

Timothy Grinsell. *Semantic Indecision*. PhD thesis, University of Chicago, Chicago, IL, 2017.

Brian Hedden. Consequentialism and collective action. *Ethics*, 130(4):530–554, 2020. doi: 10.1086/708535.

Brian Hedden and Daniel Muñoz. Dimensions of value. *Noûs*, forthcoming.

Brian Hedden and Jacob M. Nebel. Multidimensional concepts and disparate scale types. ms.

S. L. Hurley. Supervenience and the possibility of coherence. *Mind*, 94(376): 501–525, 1985.

Hans Kamp. Two theories about adjectives. In E. L. Keenan, editor, *Formal Semantics of Natural Language*. Cambridge University Press, 1975.

Rosanna Keefe. *Theories of Vagueness*. Cambridge University Press, Cambridge, 2000.

Christopher Kennedy. *Projecting the Adjective*. PhD thesis, U.C. Santa Cruz, 1997.

Christopher Kennedy. Polar opposition and the ontology of 'degrees'. *Linguistics and Philosophy*, 24:33–70, 2001.

Christopher Kennedy. Vagueness and grammar. *Linguistics and Philosophy*, 30(1): 1–45, 2007.

Christopher Kennedy and Louise McNally. Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 81(2):345–381, June 2005.

Anna B. Khmelnitskaya. Social welfare orderings for different subgroup utility scales. Discussion paper No. 198, Center for Rationality and Interactive Decision Theory at the Hebrew University of Jerusalem, 1999.

Anna B. Khmelnitskaya and John A. Weymark. Social choice with independent subgroup utility scales. *Social Choice and Welfare*, 17:739–748, 2000.

Ewan Klein. A semantics for positive and comparative adjectives. *Linguistics and Philosophy*, 4:1–45, 1980.

Thomas Kroedel and Franz Huber. Counterfactual dependence and arrow. *Noûs*, 47(3):453–466, 2013.

Andrew Y. Lee. Degrees of consciousness. *Noûs*, forthcoming. doi: 10.1111/nous.12421.

David Lewis. General semantics. *Synthese*, 22:18–67, 1970.

William MacAskill. Normative uncertainty as a voting problem. *Mind*, 125(500): 967–1004, October 2016.

Kenneth O. May. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica*, 20(4), 1952.

Sally McConnell-Ginet. *Comparative Constructions in English: A Syntactic and Semantic Analysis*. PhD thesis, University of Rochester, Rochester, NY, 1973.

Louise McNally and Isidora Stojanovic. Aesthetic adjectives. In James Young, editor, *The Semantics of Aesthetic Judgment*. Oxford University Press, 2017.

Michael Morreau. It simply does not add up: Trouble with overall similarity. *Journal of Philosophy*, 107(9):469–490, 2010.

Jacob M. Nebel. The good, the bad, and the transitivity of *Better Than*. *Noûs*, 52 (4):874–899, 2018.

Jacob M. Nebel. Utils and shmutils. *Ethics*, 131(3):571–599, 2021a.

Jacob M. Nebel. Aggregation without interpersonal comparisons of well-being. *Philosophy and Phenomenological Research*, Early View, 2021b.

Jacob M. Nebel. Ethics Without Numbers. forthcoming.

Samir Okasha. Theory choice and social choice: Kuhn versus arrow. *Mind*, 120 (477):83–115, January 2011.

Stuart Rachels. Counterexamples to the transitivity of better than. *Australasian Journal of Philosophy*, 76(1):71–83, 1998.

Galit Sassoon. A degree-approach account of multidimensional gradability. *Unpublished manuscript*, 2015.

Galit W. Sassoon. A typology of multi-dimensional adjectives. *Journal of Semantics*, 30(3):335–380, August 2013a.

Galit W. Sassoon. *Vagueness, Gradability, and Typicality: The Interpretation of Adjectives and Nouns*. Current Research in the Semantics/Pragmatics Interface. Brill, Leiden and Boston, 2013b.

Miriam Schoenfield. Chilling out on epistemic rationality: A defense of imprecise credences. *Philosophical Studies*, 158(2):197–219, 2012. doi: 10.1007/s11098-012-9886-7.

Amartya Sen. *Collective Choice and Social Welfare*. Economics. Penguin Books, third edition edition, 1970.

Amartya Sen. On weights and measures: informational constraints in social welfare analysis. *Econometrica: Journal of the Econometric Society*, 45(7):1539–1572, 1977.

Nicholas J. J. Smith. *Vagueness and Degrees of Truth*. Oxford, England: Oxford University Press, 2008.

Stephanie Solt. Multidimensionality, subjectivity, and scales: Experimental evidence. In Louise McNally Elena Castroviejo and Galit Sassoon, editors, *The Semantics of Gradability, Vagueness, and Scale Structure: Semantic Perspectives*, pages 59–91. Spring, 2018.

Larry S. Temkin. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford University Press, Oxford and New York, 2012.

Teruji Thomas. Are spectrum arguments defused by vagueness? *Australasian Journal of Philosophy*, 100(4):743–757, 2022. doi: 10.1080/00048402.2021.1920622.

Kai-Yuen Tsui and John A. Weymark. Social welfare orderings for ratio-scale measurable utilities. *Economic Theory*, 10:241–256, 1997.

Robert van Rooij. Measurement and interadjective comparisons. *Journal of Semantics*, 28(3):335–358, 2011.

John A. Weymark. Arrow's theorem with social quasi-orderings. *Public Choice*, 42(3):235–246, 1984.

Timothy Williamson. *Identity and Discrimination*. Wiley-Blackwell, 1990.

Timothy Williamson. *Vagueness*. Routledge, London and New York, 1994.