An argument for vagueness with holes

Timothy W. Grinsell

University of Chicago
Chicago, IL, USA
twg@uchicago.edu

Linguistic vagueness is a consequence of aggregating many judgments into one. For example, whether something is a *heap* depends on judgments along at least two dimensions, height and width. Results from social choice theory—a branch of economics dealing with collective decision making—show that such judgment aggregations face significant limitations. Topologically, these limitations stem from "holes" in the structure of multidimensional domains over which judgment aggregations occur. Semantically, these limitations manifest as vagueness effects like susceptibility to the sorites paradox.

1 Vagueness Distinguished

In philosophical and linguistic studies of vagueness, there is widespread agreement as to the core vagueness phenomena: participation in the sorites paradox, borderline cases, and higher-order vagueness [24, 15]. In one way or another, these phenomena reflect the difficulty of drawing a boundary between a predicates true and false applicationsbetween, for example, the *tall* things and the non-*tall* things.

The sorites paradox, also known as the paradox of the heap, is illustrated in (1).

(1) Premise 1: A pile of 10,000 grains of sand is a heap.
Premise 2: A heap of sand minus one grain is a heap.
Conclusion: A pile of 1 grain is a heap.

This paradox, most important of vagueness phenomena, is traditionally attributed to Eubulides of Miletus in the 4th century BCE. As the structure of the 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) [23].

Vague predicates also have "borderline cases." Borderline cases reflect an intuition of indecisiveness about the truth of a proposition. For example, even if we know Clarences exact height—5 feet 11 inches—we are not sure if *Clarence is tall* is true. Not only are borders hard to find between the extension and anti-extension of a predicate, but they are also hard to find between, for instance, *tall* and *definitely tall*. This is the essence of higher-order vagueness: boundarylessness in a predicates extension.[1]

[1]The problem of boundarylessness separates vagueness from ambiguity or underspecification. An ambiguous word like bank has two (or more) semantically distinct meanings; disambiguation requires solving a mapping problem between sound (or some other representation) and meaning. In contrast, vague predicates involve meanings with unclear boundaries; interpreting a vague predicate requires solving a sorting problem among the relevant entities. Thus, ambiguous words do not display vagueness effects like borderline cases, at least not on account of their distinct meanings. For example, there is generally no borderline case between "financial institution" and "shore of a river" implicated by each use of bank. 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*) [24].

2 Camping and the Sorites

I follow [23] in adopting a choice functional approach to adjectival semantics. Building on the framework of [16], van Rooij treats gradable adjectives as choice functions that select entities from the comparison class (the set of entities somehow relevant to the interpretation of the gradable adjective).[2] The gradable adjective P can be "thought of as a choice function, selecting the *best* elements of [the comparison class] c" [23, 140].

For each subset S of available options X, a *choice function* assigns to each S some element or elements of S. The rule by which the choice function decides which elements to assign is called the *choice rule*. For my purposes, then, the set of available options X will be the set of entities that constitute the comparison class.

The challenge in accounting for gradable adjectival semantics is figuring out the choice rule. For example, on one prominent theory of adjectival semantics, gradable adjectives like *tall* are measure functions that map an entity to a position on a one-dimensional scale of height [14]. These measure functions combine with a phonologically null morpheme to derive the denotation in (2).

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

Loosely paraphrased, x is tall if x's height stands out relative to some "standard of comparison," which represents the cutoff point between positive and negative extensions. In this case, the relevant choice rule is the mechanism that helps us figure out how to set the standard of comparison—how, that is, to sort the entities that fall into the positive denotation from those that don't. I will adopt this view of the choice rule in what follows.

2.1 Adjectives as Collective Choice Functions

Adjectives like *tall* are not only choice functions, they are *collective* choice functions. These adjectives aggregate multiple potential standards into a single standard.[3]

There have been multiple proposals for how the standard $\mathbf{s}(\text{tall})$ is determined. The heights of the relevant objects are important, of course, but some norm-based cut-off (like the mean) is inadequate to account for vagueness effects [11]. As Fara puts it, "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" [8]. For Fara, this includes "what our interests are." Others argue that a "stand-out" relation plays a role [15];[4] others, the distribution of

[2] 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 [16]. The set of relevant individuals is the comparison class.

[3] In some ways, this is the essential insight of supervaluationism [7]. [7] identifies two sources of judgment aggregation: verdict aggregation (as in *tall*) and dimension aggregation (as in *healthy*). One distinction between the two types of aggregators is the comparative: *taller* is not vague, while *healthier* may be. Here, I am treating adjectives like *tall* as (something like) dimension aggregators, though I leave open the possibility that verdict aggregation may be playing a role as well. One important reason for treating *healthy* and *tall* similarly is the fact that, at least on one view, the comparative *healthier* is not vague. If a possesses certain measures on all of its dimensions, and b possesses the same measures on all dimensions save one—on which b measures slightly less than a—then a is healthier than b. In this case, a small change in some value clearly demarcates true uses of *healthier* from false uses.

[4] As [15] shows, implicit comparisons like (1a) are incompatible with "crisp judgments," judgments based on small but noticeable differences in degree. In (1a), 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* (1b). In particular, both (1a) and (1b) 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, (1a), but not (1b), is infelicitous when used to make a claim about a 100-page book and a 99-page book.

heights in the comparison class [22]; others, the metric distance from a prototype [13, 6]; and still others, the shape of the probabilistic distribution of heights [17].

Contextualist, many-valued logical, and supervaluationist approaches to adjectival semantics also reflect this ambivalence. Contextualist analyses like [23], for instance, hold that the adjectival standard constantly changes. "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)" [21, 113]. Contextualist accounts place particular emphasis on the act of interpretation, arguing that the act itself changes the semantic "facts on the ground." Many-valued approaches proliferate truth values rather than standards, but the effect on interpretation mimics contextualist theories [5, 3]. And supervaluationist approaches explicitly rely on an interpretation function that simultaneously takes into account multiple adjectival standards [9, 11, 20].

Instead of choosing among these competing approaches, then, I associate *tall*'s adjectival standard with with a vector of several values—a norm-based potential cutoff, a threshold value, a probability value associated with the distribution of heights, and more. For example, *tall*'s standard may be associated with a vector in two or more dimensions, as in (3).

(3) $\mathbf{s}(\text{tall}) = (\text{height norm, ordinal rank in comparison class}, \ldots)$

I will refer to a particular vector of values $(a, b, \ldots)$ as a *cutoff*.

I assume the availability of multiple potential cutoffs in any given context. There are multiple ways to rank these cutoffs. For example, one ranking might prioritize distance from a prototype and then ordinal rank in the comparison class, while another might prioritize (positive or negative) ordinal rank in the comparison class alone. I use the term *standards* to refer to these ways of ranking. This recapitulates the central insight of supervaluationism, in which the standard-setting function looks to multiple different ways of "making precise" a vague predicate. Moreover, the availability of multiple different standards in the same context accounts for the felicity of seeming contradictions like (4a). On the present view, (4a) is really something like (4b). Contrast this with (4c), which is a true contradiction.

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

b. In one respect, Clarence is tall, and in another respect, Clarence is not tall.

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

The semantic task for adjectives like *tall*, then, is to aggregate choices among multiple standards into a single standard. This becomes the "standard of comparison." Gradable adjectives like *tall* are therefore collective choice functions, and they are subject to the limitations of collective choice.

2.2 The Limitations of Collective Choice

Social choice theory, the branch of economics dealing with the aggregation of judgments, places significant limits on whether such aggregations are possible. Arrow's Theorem represents these

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

b. This book is longer than that book.

As Kennedy 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 [(1a)] 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" [15, 18].

Figure 1: Aggregation of location preferences along a circular lake.

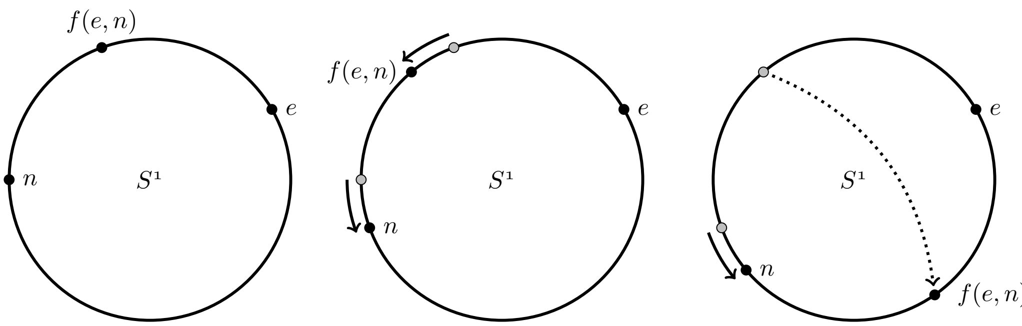

limits. It states that there is no collective decision procedure that respects certain reasonable assumptions and avoids intransitivity [1].

Generalizing Arrow's Theorem, Chichilnisky ([4]) provides a topological proof of the limits of collective choice. This proof is informally presented by [2]. 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" [4, 337]. This last requirement brings a kind of stability into the decision rule [18]. 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. 1, if Nino's preferred location n moves continuously in a counterclockwise direction, so does $f(e, n)$—until, that is, n 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 $f : S^1 \times S^1 \to S^1$ 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.

3 Vagueness Effects and the Structure of the Aggregation Domain

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

3.1 The Space of Standards Has a Hole

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 bundles of information of the type that determines potential cutoffs. Recall (a modified version of) the possible vector for the cutoff of *tall* (3).

(3) $\mathbf{s}$(tall) = (height norm, ordinal rank in comparison class)

An ordered tuple $(4, 3)$ in this context would stand for a height of 4 (in relevant units) and an ordinal rank of 3 in the comparison class. This represents a potential cutoff between the extension and anti-extension of a gradable adjective.

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. This choice rule corresponds to an adjectival standard. And different standards may rank different cutoffs differently. For example, a probability-based standard and a prototype-based standard may assign different values (and thus different relative rankings) to cutoffs $\vec{v_1}$ and $\vec{v_2}$. The diversity of adjectival standards likely means that different standards will prefer different cutoffs.[5]

Moreover, we expect the relevant standards to generate (linear) indifference sets or equivalence classes. A norm-based standard may rank the $(5, 4)$ tuple equivalently to the $(5, 3)$ tuple for the simple reason that the norm-based standard only cares about mean height, not relative position in the comparison class. The "gradient vector" of these indifference sets indicates the direction of greatest increase in along the relevant standard, as $\vec{v_{s_1}}$ and $\vec{v_{s_2}}$ in Figs. 2(a) and 2(b). In other words, the gradient vector points in the direction of cutoffs ranked highest by the relevant standard.

Finally, we consider ordinal preferences among the standards as a simplifying assumption. 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. 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 [18].

An aggregation rule is a map F (as in 5) from a tuple of standards to a single standard.

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

This recalls the camping problem (Fig. 1). 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.

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. As I show in subsection 3.3, it is the hole-y structure of the

[5] Indeed, Chichilnisky's result only works for suitably diverse voters.

Figure 2: Gradient vectors from (a) and (b) in a unit circle.

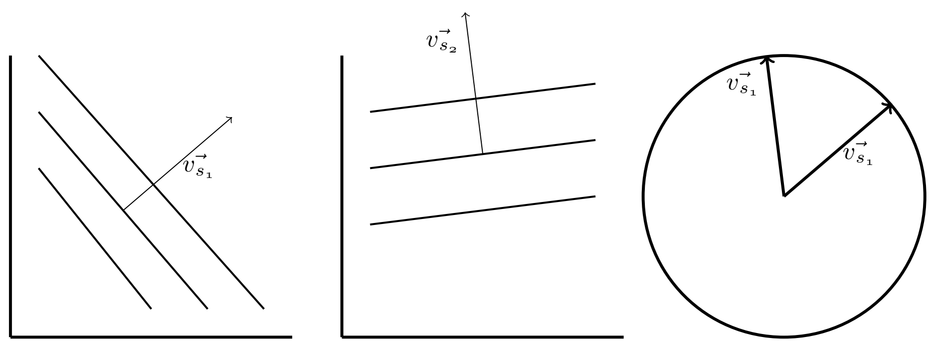

space of standards that is ultimately responsible for Chichilnisky's result and, consequently, vagueness effects.

3.2 Chichilnisky's Assumptions Are Satisfied

Accordingly, if the map F in (5) is continuous, unanimous, and anonymous, then F does not exist. Formally, the three assumptions unanimity, anonymity, and continuity look like (6).

(6) 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 $|q - p| < \delta$.[6]

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.[7] 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. Formally, the aggregation rule is invariant under permutations of the standards in its domain. At least for some uses of adjectives like *tall*, this seems like an uncontroversial assumption.

Continuity is more controversial. 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

[6] 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.

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" [10, 3].

[7] This is analagous to "supertruth" in supervaluationism [24].

to a site next to the previous one. As Lauwers ([18, 454]) explains, "Continuity is a form of stability of the social outcome with respect to small changes in the individual preferences."

Tolerance is just this kind of stability. The notion of tolerance (7) has played a prominent role in theories of vagueness.[8]

(7) a. Suppose the objects a and b are observationally indistinguishable in the respects relevant to P; then either both a and b satisfy P or neither of them does. [12]

b. $P(a) \wedge a \sim_P b \rightarrow P(b)$

Vague predicates are said to be tolerant, as the second premise of the sorites paradox (8) demonstrates.

(8) Premise 1: A 2m tall woman is tall.
Premise 2: Any woman 1mm shorter than a tall woman is tall.
Conclusion: Therefore, a 1m tall woman is tall.

A vague predicate like *tall* is tolerant in the sense that small differences in height seemingly do not change the value of $[\![\text{tall(woman)}]\!]$ from 1 to 0. Put another way, as the heights of x and y converge, the values $[\![\text{tall}(x)]\!]$ and $[\![\text{tall}(y)]\!]$ also converge. This implies that a small change in height should not be the difference between the *tall* entities and the non-*tall* entities. That is, vague predicates are "stable" in the manner demanded by continuity. Thus, the principle of tolerance—a key property of the sorites premise—approximates the work of continuity.[9]

When defined in terms like (6c), continuity emphasizes what Saari ([19, 57]) calls "local behavior": if x_o and x_n are so close together that they are difficult to distinguish, so must be $f(x_o)$ and $f(x_n)$ (where f is a continuous function). But this ignores the global structure of the domain, and when the local structure and the global structure disagree, problems arise. Saari ([19, 57]) provides the following example: "If we lived on a huge circle and used only local information, we would view the circle as being, essentially, a straight line." And the camping problem demonstrates the perils of treating circles as lines.

3.3 Vagueness with Holes

If the domain of the aggregation function contains holes, then unanimity, anonymity, and continuity are mutually incompatible. Vagueness effects result from this incompatibility. For example, the sorites paradox represents a tradeoff between conflicting assumptions. The first premise and the conclusion reflect the assumption of unanimity: every standard agrees that a 2m tall woman is *tall* and that a 1m tall woman is not *tall*. The sorites premise reflects continuity and anonymity. According to continuity, a small change in a tall woman's height should not take her from *tall* to not *tall*. According to anonymity, no contextually relevant standard is somehow more important than another. The sorites paradox is paradoxical because it rests on inconsistent assumptions.[10]

[8] See [5] for a recent tolerance-based approach.

[9] Despite this affinity between tolerance and continuity, [21] argues that continuity is too strong a constraint to model tolerance. First, continuity overgenerates the number of vague predicates. Second, sometimes domains are discrete. For example, the sorites premise for *heap* works with grains of sand, not fractional grains of sand. Topologically speaking, every discrete domain could support a continuous function, but we want to distinguish between "discrete" vague predicates like *heap* and "discrete" non-vague predicates like *has 100 hairs or fewer.*
Instead, Smith ([21]) 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.

[10] Borderline cases result when the aggregation function fails to return a value for non-unanimous aggregations.

Figure 3: Loops on a closed disk and an annulus.

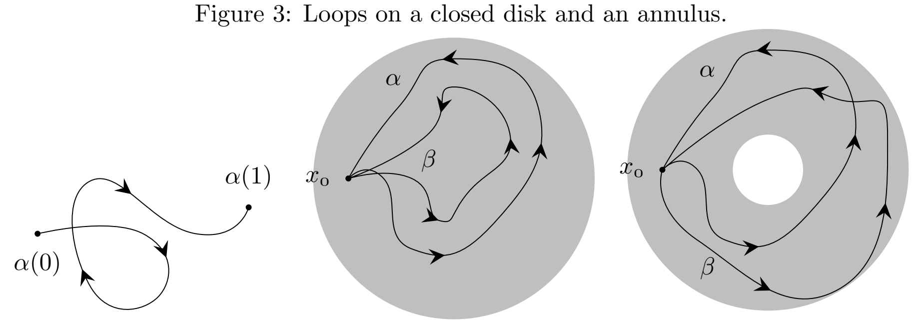

(a) (b) (c)

This inconsistency only arises in choice spaces with holes. To see what I mean, 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.

To prove this, we need the notion of a path.[11] 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. 3(a)). A "loop" is a path in which $\alpha(0) = \alpha(1)$. Usually, this point is identified at x_o, as in Fig. 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.

In Fig. 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_o and goes nowhere in between (that is, the constant loop is the point x_o). And we can do this for all loops in the closed disk. Therefore, the set of equivalence classes of loops starting at x_o in the closed disk is just the set of equivalence classes of the constant loop.

However, consider the annulus in Fig. 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. 3(c)). One can walk along the path α one time, or two times, or three .... And one can walk along

[11] I rely heavily on [18] and [2] in what follows. See these references for a more detailed account of Chichilnisky's proof.

the path α in the opposite direction once, twice, etc.

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. 3(c).

We can now prove (a simple version of) Chichilnisky's theorem and its linguistic corollary. We associate the set of integers $\mathbb{Z}$ with the equivalence classes of loops in S^1 (once around is 1, twice around is 2, etc.) and the set of ordered pairs of integers $\mathbb{Z} \times \mathbb{Z}$ with loops in $S^1 \times S^1$. The loop in which one individual standard makes a complete rotation but a second individual standard does not is $(1,0) \in \mathbb{Z} \times \mathbb{Z}$. The opposite loop (in which the second individual standard makes a complete rotation but the first does not) is $(0,1)$. Sequential positive rotations are thus represented by $(1,0)+(0,1)$. Simultaneous positive rotations are represented by $(1,1)$.

In this integer-oriented framework, an aggregation function is a function $F_* : \mathbb{Z} \times \mathbb{Z} \to \mathbb{Z}$. Anonymity implies (9) because it doesn't matter which standard does the rotating (the aggregated standard depends on total rotations, not how the total is distributed among the standards).

(9) $F_*((1,0)+(0,1)) = F_*(2,0) = 2F_*(1,0)$

But $(1,0)+(0,1) = (1,1)$. And since unanimity implies that $F_*(1,1) = 1$, we have (10).

(10) $1 = F_*(1,1) = 2F_*(1,0)$

This means that there must be some integer $F_*(1,0)$ such that 2 times this integer equals 1. Impossible.

Vagueness effects like the sorites paradox are thus a product of constrained choice over spaces with holes. Evidence for this approach comes from non-vague adjectives, so-called "absolute" adjectives like *full*. Absolute adjectives do not easily give rise to the sorites paradox because the second premise is judged to be false (11).

(11) # Premise 2. Any theater with one fewer occupied seat than a full theater is full.

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*.

(12) **s**(full) = (volume)

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 [4, 347].

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 $\mathcal{P}$.)

The property distinguishing absolute from adjectives like *tall* 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. In particular, multidimensionality creates holes.

References

[1] Kenneth J Arrow. A difficulty in the concept of social welfare. *The Journal of Political Economy*, pages 328–346, 1950.

[2] Nicholas Baigent. Topological theories of social choice. In Kenneth J Arrow, Amartya Sen, and Kotaro Suzumura, editors, *Handbook of social choice and welfare*, volume 2, pages 301–334. Elsevier, 2010.

[3] Heather Burnett. A delineation solution to the puzzles of absolute adjectives. *Linguistics and Philosophy*, 37(1):1–39, 2014.

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

[5] Pablo Cobreros, Paul Egré, David Ripley, and Robert van Rooij. Tolerant, classical, strict. *Journal of Philosophical Logic*, 41(2):347–385, 2012.

[6] Igor Douven, Lieven Decock, Richard Dietz, and Paul Égré. Vagueness: A conceptual spaces approach. *Journal of Philosophical Logic*, 42(1):137–160, 2013.

[7] Paul Égré, Christian List, and Franz Dietrich. Supervaluationism: a judgment-aggregation perspective. Presentation at VII Navarra Workshop on Vagueness, December 2013.

[8] Delia Graff Fara. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics*, 28(1):45–82, 2002. Originally published under the name Delia Graff.

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

[10] Geoffrey Heal. Social choice and resource allocation: a topological perspective. In Geoffrey Heal, editor, *Topological social choice*, pages 1–14. Springer, 1997.

[11] Hans Kamp. Two theories about adjectives. In Edward Keenan, editor, *Formal semantics of natural language*, pages 123–155. Cambridge University Press, 1975.

[12] Hans Kamp. A theory of truth and semantic representation. In Paul Portner and Barbara H. Partee, editors, *Formal semantics: the essential readings*, pages 189–222. Blackwell, 1981.

[13] Hans Kamp and Barbara Partee. Prototype theory and compositionality. *Cognition*, 57(2):129–191, 1995.

[14] Christopher Kennedy. *Projecting the adjective: The syntax and semantics of gradability and comparison*. Routledge, 1999.

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

[16] Ewan Klein. A semantics for positive and comparative adjectives. *Linguistics and philosophy*, 4(1):1–45, 1980.

[17] Daniel Lassiter and Noah D Goodman. Adjectival vagueness in a bayesian model of interpretation. *Synthese*, pages 1–36, 2015.

[18] Luc Lauwers. The topological approach to the aggregation of preferences. *Social Choice and Welfare*, 33(3):449–476, 2009.

[19] Donald G. Saari. *Disposing dictators, demystifying voting paradoxes*. Cambridge University Press, 2008.

[20] Galit W. Sassoon. *Vagueness, Gradability and Typicality: The Interpretation of Adjectives and Nouns*. Brill, 2013.

[21] Nicholas JJ Smith. *Vagueness and degrees of truth*. Oxford University Press, 2008.

[22] Stephanie Solt and Nicole Gotzner. Experimenting with degree. In Anca Chereches, editor, *Semantics and Linguistic Theory*, volume 22, pages 166–187, 2012.

[23] R. van Rooij. Vagueness and linguistics. In Giuseppina Ronzitti, editor, *Vagueness: A guide*, pages 123–170. Springer, 2011.

[24] Timothy Williamson. *Vagueness*. Routledge, 1994.