BOUNDARIES AND PRUDENCE: WHY YOU DON'T CARE ABOUT

THE SORITES PARADOX

by

ERIC PAUL SNYDER

(Under the Direction of Yuri Balashov)

ABSTRACT

The following concerns vagueness and the Sorites Paradox. It attempts to deflate the significance of the paradox via broadly 'pragmatic' considerations. First, I argue that vagueness is a necessary feature of natural languages, i.e. we could not do without vague expressions. Second, I give necessary conditions for Sorites construction and argue that its plausibility relies on our presupposing a certain metalinguistic imperative, itself deriving from the necessity of vagueness. Finally, I give the prevailing semantics for gradable predicates and show that, in general, if this semantics is correct, then the Sorites poses no threat to the semantics of prototypical vague predicates. If my arguments are on track, we have a nice explanation for why this paradox has remained obstinate for so long: We are searching for 'hidden' boundaries that do not and, in fact, could not exist. Rather, the problem is essentially with classical logic and set-theory.

INDEX WORDS: Vagueness, Sorites Paradox, Contextualism, Pragmatic Presupposition, Gradable Adjectives

BOUNDARIES AND PRUDENCE: WHY YOU DON'T CARE ABOUT

THE SORITES PARADOX

by

ERIC PAUL SNYDER

B.A., Piedmont College 2005

A Thesis Submitted to the Graduate Faculty of The University of Georgia in Partial Fulfillment

of the Requirements for the Degree

MASTER OF ARTS

ATHENS, GEORGIA

2009

BOUNDARIES AND PRUDENCE: WHY YOU DON'T CARE ABOUT

THE SORITES PARADOX

by

ERIC PAUL SNYDER

Major Professor: Yuri Balashov

Committee: Charles Cross
René Jagnow

Electronic Version Approved:

Maureen Grasso
Dean of the Graduate School
The University of Georgia
August 2009

ACKNOWLEDGEMENTS

I would like to thank Yuri Balashov, Charles Cross, and René Jagnow for their help and patience in regards to this thesis, classes, and conversations. Special thanks are due to Yuri, for whom this thesis would not have been possible without.

TABLE OF CONTENTS

Page

Section 1: Introduction

Traditional theories of Vagueness and the Sorites paradox fall into three categories, those based on revisions to classical logic, e.g. multi-valued logics or infinite-valued logics, those based on revisions to classical semantics, e.g. supervaluationism, and those who see the problems posed by vagueness as essentially epistemic in nature. More recently, "pragmatic" (very broadly understood) accounts have gained increasing support. Contextualists, for instance, think features of conversational context are crucial to understanding vague expressions. Others emphasize how the assertibility conditions of sentences containing vague expressions differ from those of other sentences. Owing to the truly enormous literature and variety of views, a comprehensive treatment of the topic would be futile for the purposes of a master's thesis. Instead, I propose to narrow the lens to contextualist theories, which I take to be the best general strategy currently available. Again, however, owing to the variety of positions within this general approach, I cannot outline in any great detail each of the viable positions. Rather, I will touch on several contextualist theories throughout, giving the general positions and relevant distinguishing features.

The program I will present is ambitious. Apart from outlining current contextualist theories of Vagueness and solutions to the Sorites Paradox, I will focus on certain important features of vagueness in natural languages and aspects of the Sorites paradox which have thus far been largely neglected within the literature. These include:

(i) game-theoretic, evolutionary considerations for the usefulness of vague expressions

(ii) (as a corollary) the necessity of vagueness within natural languages

(iii) certain pragmatic considerations on the assertibility of sentences containing vague expressions

(iv) constraints necessary for legitimate Sorites paradox construction

(v) the fact that not all Sorites paradoxes are equally plausible

(vi) (as a corollary) that Sorites paradoxes themselves are susceptible to meta-level Sorites paradoxes

(vii) the semantics of prototypical vague predicates like 'tall', 'bald', and 'heap'

(viii) the fact that the semantics of prototypical vague predicates does not transition smoothly to Sorites paradoxes

(ix) a 'normative vs. descriptive' distinction inherent in the Sorites, seemingly needed for a cogent explanation of its infamous obstinacy

Owing to the breadth of (i) – (ix) and considerations of brevity, much of the exposition may seem somewhat hurried and in need of further detailed support. A significant reason for this, I think, is owing to the novelty of the project. First, my account does not fit nicely within any of the aforementioned general positions. This has the advantage of being unpartisan, meaning that no prominent theories are necessarily precluded, and theorists from all camps can benefit from any insights. Another advantage is that I can view vagueness and the Sorites without any particular theoretical prejudices. In a sense, this necessitates the breadth of scope in (i) – (ix). Secondly, it does not directly address the problems posed to classical logic by the Sorites, which have commonly taken precedence in traditional theories of Vagueness. This is simply a consequence of the aforementioned. The benefit of bringing to the forefront certain traditionally neglected but important aspects of vagueness and the Sorites, I hope, outweighs the cost of insufficient detail.

The thesis is divided into three parts, each containing subsections. In the first part, I give a general exposition of vagueness as it is understood in the contemporary literature. I do this with a mind to fairness to competing theories – something more difficult than might initially be thought. I also offer a thought experiment aimed at showing the necessity of vagueness in natural languages. I do this in terms common to evolutionary game-theory and its application to natural languages. Finally, I emphasize the usefulness of vague expressions. Using evidence cited in recent work in pragmatics, I argue that the purpose of vagueness in natural languages is to supply a sufficient amount of semantic information without being overly precise. Being overly precise is, in effect, inefficient. In the second section I argue two theses. The first is that not all versions of the Sorites paradox are equally plausible. The second is that we do not take any inductive premises of the Sorites for granted. I appeal to (broadly) Gricean principles governing the assertibility of sentences containing vague expressions. I suggest that an appropriate analysis of vagueness should appeal to the assumptions, presuppositions, and other propositional attitudes operative in conversational contexts, or what Robert Stalnaker calls *common ground*. The third part is intended to unify considerations in the first two parts and to apply them to the Sorites. First, I sketch the commonly accepted semantics for gradable adjectives, the grammatical category of commonly cited vague predicates. The result is that these rely intimately on features of conversational context, and that there are no objective, non-context-sensitive properties of baldness, tallness, *etc*. After this, I examine Delia Graff Fara's contextualist account and its problems. Ultimately, I suggest that the evolutionary considerations and appeals to communicative efficiency are needed to generalize her account, which I take to be the best currently available.

The upshot of the arguments I give is a deflationary view of the Sorites. That this ancient paradox has plagued philosophers and others for a couple millennia now needs explaining. As a tentative hypothesis, it may very well turn out that the semantics of our ordinary vague expressions has not, and I will argue *could* not, evolved in such a way to be able to carve reality with any arbitrary degree of precision. Since this is exactly what the Sorites demands of these expressions, we have good *prima facie* reason for thinking that searching for arbitrarily precise boundaries is futile and, hence, that our fascination with the paradox is perhaps ill-founded. At a minimum, we should reevaluate our interests and expectations in entertaining the paradox. My hope is that this thesis motivates such a reevaluation.

Section 2: Vagueness in Natural Languages

What is Vagueness?

This question is a bit like that question asked on the first day of any introductory philosophy class – "What is Philosophy?" – in that an attempt to answer the question presupposes the work needed to be done, and that whatever answer is given is likely to offend another's account. To say what 'philosophy' is requires engaging in philosophy, and one's conception of philosophy will inevitably conflict with another's. Similarly, to say exactly what 'vagueness' is presupposes that a somewhat concise characterization can be given, and this characterization is likely to preclude other accounts. To complicate matters further, just as what is to count as 'philosophical' requires some kind of philosophical argument, just what is to count as 'vague' is itself vague. Nevertheless, philosophers and linguists have insisted on there being an interesting distinction between 'vagueness' as it is ordinarily used and 'vagueness' as a peculiar semantic property of natural language with certain identifiable features. Some kind of characterization is needed.

There are a few options. We might introduce some prototypically vague predicates such as 'bald', 'heap', or 'tall' and say that vagueness is the semantic property that these and similar words share. The problem here, of course, is deciding exactly what 'similar' means, and this is just to ask our initial question again. Another problem is that predicates are not the only vague lexical items: Many singular terms, adverbs, quantifiers, and modifiers (at least) are commonly considered vague as well. Alternatively, we might say that vagueness is exemplified by 'borderline cases' of application. A *borderline case* of a predicate *F* is an object for which it is unclear whether *F* definitely applies or definitely does not apply. On this idea, an expression is vague if it has borderline cases of application. The problem here is that we would like to know what 'unclear' and 'definitely' mean, but to give anything other than their ordinary, intuitive meanings would be to unfairly rule out accounts which crucially understand these locutions differently.

Similarly, we might say that an expression is vague if it lacks a *well-defined extension* (contains *fuzzy boundaries*). A predicate lacking a well-defined extension is one lacking a sharp boundary between its positive and negative extensions. Since this just means that there are some objects which clearly satisfy the predicate, some which clearly do not, and an intermediate range which neither clearly satisfy nor clearly do not satisfy the predicate, it is not obvious this is really an improvement over the previous suggestion. Furthermore, this characterization unfairly militates against a particular view (epistemicism) holding that all vague expressions actually do have well-defined extensions, it's just that we're epistemically cut off from them.

Finally, we might say an expression is vague if it is susceptible to Sorites paradoxes. We will discuss the Sorites in more detail later, but for now it will do to say the following: For any Sorites-susceptible predicate, it must be capable of partitioning objects so that they form a natural continuum, forming a total linear ordering, where for any adjacent pair in the ordering, there is no relevant significant difference between those objects in respect to the predicate. The problem here is that "significantly different" is itself vague and, as we've already said, predicates are not the only vague expressions. Also, it's been claimed that Sorites susceptibility is not a necessary condition for vagueness, i.e. there are vague predicates which do not form natural continuums required for legitimate Sorites construction.[1]

[1] For example, Soames (1999) claims 'vehicle' and 'machine' are vague predicates that do not form natural continuums along any obvious dimension.

Clearly, it's difficult to give an adequate characterization of vagueness without some kind of circularity or question-begging. To compound difficulties, it is widely acknowledged that vagueness is pervasive. Some have suggested that *every* non-mathematical predicate suffers from some degree of vagueness. If so, then *any* characterization we are going to give will itself be vague. Let's just accept that characterizing vagueness precisely is ultimately futile and that an adequate enough description is that vague expressions more or less share all the above features. Moreover, like most theorists we will narrow our inquiry to just vague predicates. Following Crispin Wright (1975), we agree that the most prominent feature of vague predicates is that they are "tolerant". This means that "there is a notion of degree too small to make any difference" to their application.[2] This unifies the absence of "significant difference" required for the Sorites with vague predicates lacking *sharp boundaries* (epistemicism notwithstanding), i.e. their lacking well-defined extensions. 'Borderline cases' are then those objects which are not clearly in the positive extension or the negative extension of the predicate. Put another way, it is indeterminate (in the ordinary sense) whether a vague predicate applies or fails to apply to a borderline object. Prototypical vague predicates like 'bald', 'heap', and 'tall' are simply convenient heuristics for pointing out these features. I assume this characterization is intuitive enough to make progress.

Humans, Superhumans, and the Evolution of English

[2] Wright (1975, pp. 333)

As a variation on a thought experiment given by Timothy Williamson, consider the following.[3] Suppose there were two types of humans, normal humans and superhumans. Superhumans are like normal humans except for their having infallible perceptual faculties. If there are *x* grains in a pile, they know this. If there are *y* hairs arranged in such and such a fashion, they know this. If a person is exactly *z* millimeters in height at point <*x*,*y*,*z*> in a precise region of space *R* at time *t*, they know this. *Etc.* The relevant question here is: Amongst the superhumans, is there vagueness? And if so, do they speak vaguely? In one sense, they have no need for vagueness. Instead of using 'heap', 'bald', or 'tall', they might say something like the following instead:

(1) *S* has 21,342 hairs of average diameter *x* arranged *in such and such a manner* on *S*'s head

(2) Pile A contains 14,234 grains of salt of average diameter *x* arranged *in such and such a manner*

(3) *S* is exactly 2.546… meters in height at coordinate <*x*,*y*,*z*> in a precise region of space *R* at time *t*

On the other hand, the superhumans may not care to speak so scrupulously. After all, speaking this way would take a long time. Moreover, there do not appear to be any ordinary communicative situations in which this kind of precision would be crucial to their overall well-being. They need to eat, sleep, *etc*, and the (potentially very long) time it would take to always speak with this exactness would militate against these needs. So it appears the superhumans have a *prima facie* reason for speaking vaguely (for ordinary purposes).

We can put the point in terms of evolutionary game theory. In evolutionary game-theory, strategies are considered the 'players' so that strategies themselves are modeled as playing against one another in recurring dynamic games. One strategy is considered better than another if

[3] Williamson (1994, pp. 198-204) considers a thought experiment with omniscient speakers undergoing the 'dynamic' version of the Sorites.

it is more likely to leave more copies of itself in the next generation, i.e. when the next 'game' is played. A strategy is said to be at an *Evolutionary Stable Strategy* (ESS) equilibrium if no individual playing that strategy could improve its reproductive fitness by changing to another strategy, and if no 'mutant' playing a different strategy could 'invade', or establish itself in, the population.[4] It is easy to see that speaking vaguely is an ESS equilibrium for ordinary humans – since they cannot speak with the precision of superhumans, this cannot be a better strategy for them, nor can superhuman precision invade their population for the same reason. It is less obvious that speaking vaguely is a superhuman ESS equilibrium.

Without assigning arbitrary utility functions, it seems undeniable that the superhumans will value their time for meeting everyday needs over the (potentially long) time spent speaking with the kind of precision described in (1) – (3). Even if there are times in which it is beneficial to speak precisely, surely these will be far fewer than the alternative. After all, it is difficult to imagine a situation in which speaking precisely would ordinarily offer any real advantage over speaking imprecisely. This is because the imprecise correlates to (1) – (3) convey enough required information on their own. For example, nothing of obvious import ordinarily turns on exactly how many hairs one has, and if a superhuman were interested in conveying that someone has less that average hair on their scalp, 'bald' would be sufficient. In this sense, imprecision would be more communicatively efficient than precision even if less informative – the information lacking simply is not important for their ordinary purposes. Since it is reasonable to assume that communicative efficiency is the most prominent factor determining the linguistic behavior of the superhumans, we have good reason to conclude that speaking vaguely is indeed an ESS equilibrium for them as well as the ordinary humans.

[4] See Smith (1982).

How, if at all, does any of this carry over to the evolution of natural languages? It is natural to think that words exist in a natural language in as much as they help facilitate effective communication. Moreover, it is natural to think that words persist in natural languages according to their usefulness to their users. Words enter into and exit out of natural languages for certain reasons. The best explanation of this fact is that the interests and purposes of humans change over time, and certain words best fit these interests and purposes according to their usefulness. Theorists within evolutionary linguistics take this for granted. For example, both Clark (1997) and Tomasello (2003) assume that the vocabulary of natural languages evolves to suit its users' communicative needs. They also assume that natural language syntax evolves so as to suit the pre-existing cognitive and processing capacities of its users. In particular, they find it reasonable to expect languages to evolve according to the pre-existing psychological, perceptual, and motor faculties of humans developing those languages. Furthermore, some evolutionary linguists have found it reasonable to assume that these faculties constraining the evolution of languages will have been selected for language-specific tasks. As a result, both natural languages and those capacities responsible for its evolution will have coevolved.

If this coevolution thesis is correct, then vagueness regarding observational predicates (to be explained shortly) is a necessary feature of natural languages. Our faculties are not infallible, so it is not surprising that our language has evolved to reflect this fact. That is, since we do not have superhuman perceptual faculties, it should not be surprising if in fact our ordinary vague predicates do not denote some precise number of hairs, millimeters, or salt grains. For if they did, we would expect a few things. First, we would expect humans to have the capacity to casually make such discriminations. Because of this, we would (secondly) expect the actual facts (the exact number of hairs, millimeters, or salt grains) and linguistic behavior (classifying objects as ‘bald’, ‘tall’, or ‘heaps’) to be rather homogeneous. Third, and most importantly, since humans seem to have access to the meanings of their terms, we should expect them to be to identify at least broadly correct values for the variables in sentences like (1) – (3). However, none of these expectations is generally fulfilled. Surely our perceptual fallibility is crucial to explaining why this is the case. However, even if the coevolution thesis turns out to be false, we still have no better explanation for the usefulness of vagueness. It is to this theme we now turn.

Vagueness is Useful

Theorists of vagueness have often overlooked its utility. This is not shocking, given it is rather obvious. When mentioned, it is usually only in contrast with the now antiquated Fregean / Russellian view of vagueness as a “defect” of natural languages. For example, Williamson says the following: “[V]agueness is a desirable feature of natural languages. Vague words often suffice for the purpose in hand, and too much precision can lead to timewasting and inflexibility.”[5] In is a similar vein, Wright says “the utility and point of the classifications expressed by many vague predicates would be frustrated if we supplied them with sharp boundaries.” [6] Instead of being simply a paltry point aimed at an antiquated attitude, I submit, this seemingly insignificant piece of information (i.e. that vagueness is useful) is crucial to understanding the obstinacy of the Sorites paradox. So we do well to give it a little more attention.

We saw in the last section that the usefulness of speaking vaguely gave us good reason for thinking that even if we could speak without vagueness, we would be better off speaking

---

[5] Williamson (1994, pp. 70-71).
[6] ibid pp. 330.

vaguely. Recent linguistics research supports this claim. Jucker, Smith and Lüdge (2003) identify several purposes for speaking vaguely. To summarize their findings:

> Vagueness is not only an inherent feature of natural language but also – and crucially – it is an interactional strategy. Speakers are faced with a number of communicative tasks, and they are vague for strategic reasons. Varying the level of vagueness may help guide the addressee to make the intended representation of entities and events and to draw intended implications from them…[S]peakers constantly negotiate their common ground, seeking and providing cues as to the partner's beliefs and the current accessibility of beliefs that are relevant to the interpretation of an utterance…[V]ague expressions are not just poor but good-enough substitutes for precise expressions, but are preferable to precise expressions because of their greater efficiency.[7]

This "greater efficiency" is exactly what we would expect under the assumption that vague expressions ordinarily convey enough information, and that whatever missing information that would have been communicated by a precise correlate is generally unnecessary given speakers' ordinary intentions.

It is commonly recognized that prototypically vague predicates are *observational* in the sense that judgments about their applicability to objects are based on ostension and subsequent extrapolation. We learn the meaning of 'bald' by a limited number of ostensive instances, and we judge further instances by extrapolating from that initial learning. A few things follow. First, since language users will have different ostensive training, we should expect speakers to differ over their judgments concerning classifications of objects via vague predicates. Second, since the justification of speakers' judgments will essentially differ, speakers will have no obvious recourse for definitively settling classificatory disputes. Initially, this might seem like quite a problem. If we cannot help but differ over many of our classificatory judgments, and we have no principled way of arbitrating those disputes, how can communicating with vagueness be effective after all? Using game-theoretic considerations, Parikh (1996) has shown that vague

[7] Jucker, Smith and Lüdge (2003, pp. 1739).

concepts, although generally non-isomorphic between individuals, still substantially decrease processing effort. For example, even if two individuals differ in their judgments concerning which objects count as “blue”, that one can guess with regular success what the other will probably judge “blue” substantially decreases search time by ruling out other probable “non-blue” objects.

This is to be expected under the assumption that it is within the best interest of the linguistic community to keep as close to regular consistency in their usage as possible. Even if you have good reason to expect your concept “blue” to be somewhat different from mine, you do better to assume that my concept will be close enough to yours so that for the majority of objects we will agree on whether or not they are blue. Indeed, we should expect disagreements only for those objects closer to the borderline of the predicate. Without this assumption, we have no real ground for judging linguistic competency. For suppose *S* regularly judges a clear case of blue to be not blue. Without the assumption that use of ‘blue’ is somewhat regular across the linguistic community, it would be futile to appeal to linguistic convention in correcting *S*. We could give *S* a color chart and show *S* that this particular object matches ‘blue’ on the chart; but since it is linguistic convention which determines which predicate is associated with which color on the chart, and since ‘linguistic convention’ entails regularity within the linguistic community, our appeal would be circular. On the other hand, we should expect regular disputes over borderline cases since, typically, ostensive definitions are given by prototypes or exemplars, and to be a borderline case in some sense means to be undecidable between competing exemplars.

So here is the picture so far. We learn vague predicates ostensively and extrapolate on this basis to further future objects. Although speakers will typically differ over borderline cases, linguistic efficiency ensures that classifications will generally coincide across the population.

Two mutually consistent theses explain this. First, the coevolution thesis is correct. That is, human capacities and natural languages coevolved. If so, then vagueness is efficient by default, given the usual assumption that efficiency determines fitness and that vague observational predicates are the linguistic manifestation of our perceptual fallibility. Alternatively, vagueness is efficient because it conveys enough relevant information in less time and with less effort than precision for ordinary conversational purposes. We could in principle count exact number of hairs, millimeters of height, grains of salt, *etc*, but for ordinary purposes nothing crucial turns on this. The best explanation for this is that our perceptual faculties are limited, and we should expect natural languages to evolve to suit their users' limitations. People create their languages; why would they create them to be of little or no use?

Section 3: Common Ground and the Sorites

The Sorites

There are many forms of this well known paradox. Perhaps the most common is the *conditional Sorites*. For some predicate F and ordered sequence of objects $<o_1,\ldots,o_n>$[8],

(CS) P1: Fo_1
P2: $\neg Fo_n$
IP: $(\forall o)(Fo_i \rightarrow Fo_{i+1})$
C: Fo_n

where the following conditions are met:

(S1) F must be capable of partitioning $<o_1,\ldots,o_n>$ according to their degree of '*F*-ness' along some single, specifiable dimension (e.g. 'number of hairs' or 'millimeters in height')

(S2) This partitioning must create a total linear ordering of objects according to their '*F*-ness', i.e. the ordering relation on objects is antisymmetric, transitive, and total[9]

(S3) The partitioning must be fine enough to guarantee that there is no ordinarily discernable difference between adjacent units, i.e. for any adjacent pair in the ordering, there is no (casual) observable difference between those objects along the relevant dimension

(S4) The first object in the series must be definitely F, the last definitely not-F, with a suitable range of definite F's succeeding the first and a suitable range of definite non-F's preceding the last

'Bald', 'tall', and 'heap' are common Sorites predicates. Partitioning the domain of men according to number of head hairs, beginning with zero and serially increasing to an arbitrarily large number, satisfies S1–S4. S1 is met since the men are grouped along a specified dimension,

---

[8] A note on convention. Since objects are assumed to already meet S1 and S2, read IP as saying "for all objects in the domain, if the first is F, then so is the one immediately following it".

[9] A relation is antisymmetric if for any x and y, if Rxy and Ryx, then $x = y$. A relation is transitive if for any x, y, and z, if Rxy and Ryz, then Rxz. A relation is total if for all x and y, either Rxy or Ryx. In essence, a Sorites series / continuum mirrors the natural numbers ordered by '$\leq$'.

namely numbers of head hairs. S2 is met because all men are ordered from least amount of hairs to greatest amount of hairs. S3 is met because we cannot ordinarily discern a difference of one hair. Finally, S4 is met because a man with zero head hairs is obviously bald, a man with an arbitrarily large number of head hairs is obviously not bald, and because there will be many men immediately succeeding the first which are bald and many men preceding the last which are not bald. Similar comments hold for 'tall' partitioned by millimeters and 'heap' partitioned by grains of sand. The paradox results from our seeming inability to deny P1, P2, or IP. We reach C simply by *n* applications of universal instantiation and *modus ponens*. Something here must be given up, and most theorists claim it is IP.

All of the constraints are necessary for Sorites construction. S1 delivers a range of objects according to their satisfaction or non-satisfaction of the relevant predicate. S2 delivers a natural continuum of objects leading from definite *F*'s to definite non-*F*'s. S3 is necessary for the plausibility of the inductive premise. Importantly, if there were discernibly large enough degrees of *F*-ness between adjacent items, we might actually be justified in drawing a sharp boundary between the *F*'s and non-*F*'s. Finally, S4 guarantees that P1 and P2 of the paradox are beyond doubt. Furthermore, the range of definite *F*'s and definite non-*F*'s covaries with the "fineness" determined by S3 and, crucially, just what is to count as an "ordinarily discernible difference" is itself vague, resulting in *higher-order vagueness*.

Of the four constraints, S3 is most significant. It guarantees that the predicate in question is in fact tolerant. Again, for Wright, a predicate is *tolerant* if, for any predicate *F* and adjacent objects in the series o_i and o_j, "there is a notion of degree too small to make any difference" to the application of *F* to o_i and o_j.[10] Moreover, as noted, 'bald', 'tall', and 'heap' are *observational predicates* in the sense that judgments about their applicability to objects are wholly based on

[10] Wright (1975, pp. 333).

casual observation. Wright says that "The information of one or more senses is decisive of the applicability of an observational concept; so a distinction exemplified in a pair of sensorily-equivalent items cannot be expressed by means solely of predicates of observation, for any observation expression applying to either item must apply to both." [11] Wright's comments suggest the following principle:

(WP) For any tolerant predicate *F* and objects o_i and o_j, if o_i and o_j are observationally indistinguishable or saliently similar (along some relevant dimension)[12], then Fo_i if, and only if, Fo_j.

It is specifically this principle, I will argue, that is responsible for the plausibility of any legitimate Sorites inductive premise.

The Contextualist Solution

Suppose that we led a subject from the first object in the series to last by presenting every object individually and in serial order. By construction, the first object and a suitable number following it will be undeniably *F* – the subject has no choice but to judge them as *F*. At some point in the series the subject will feel less confident in the objects' decreasing *F*-ness. Assuming that he must judge each object as only *F* or not-*F*, he will continue judging objects *F* even for those borderline cases he might normally judge not-*F* (e.g. if he were looking at the spectrum as a whole). Eventually, he will realize the objects are becoming more and more not-*F* and will "jump" to calling subsequent objects not-*F*.[13] After making this transition, he might realize that

---

[11] ibid pp. 338.

[12] Assuming ordinary perceptual conditions, e.g. under ordinary lighting conditions, not wearing colored glasses, *etc*.

[13] Stewart Shapiro calls the transition from judgments of *F*'s to non-*F*'s a "jump". See his 2006 pp. 24-36 for a detailed discussion of the dynamic Sorites. In this version of the paradox, we imagine presenting individual objects

some of the immediately preceding objects are more legitimately non-*F* than *F*. If allowed, he will then begin reneging his immediately previous judgments. He will probably even judge those borderline cases he would normally judge *F* to be not-*F*. Just as before, he will reach a point at which objects appear more *F* than not-*F* and make another jump. Depending on his patience with us, he might continue in this fashion, sliding back and forth, for some time. Eventually, because of exhaustion or frustration, he will stop sliding and simply pronounce that this one is *F* and this next one not-*F*, thereby drawing a boundary. Assuming that this boundary is established somewhere near the center of the spectrum of objects, his decision will appear completely arbitrary, for a number of surrounding pairs would have been equally qualified to mark the boundary.

This is called the "dynamic" version of the paradox. Contextualists emphasize that the reason our subject will judge those borderline objects he would normally judge not-*F* to be *F* (and *vice versa*) is owing to his desire to remain consistent with prior conversational commitments. In general, we try not to contradict ourselves. Intuitively, since the objects are ordered such that the difference in the degree of *F*-ness between any adjacent pair is practicably negligible, judging some object as *F* and the next as not-*F* would (for ordinary purposes) be contradictory. Contextualists such as Scott Soames add that each new judgment sets a new conversational standard, and this standard necessitates that we judge like objects alike.[14] This explains how someone could be led from one end of the spectrum to the next judging them all (not) *F*, the so-called "forced march" version of the paradox.[15] This also provides a neat solution

---

in the series to participants in serial order. The subjects are required to classify each object as *F* or not-*F*. Beginning with the first object and proceeding all the way to last (or *vice versa*) is called the "forced marched" version of the paradox, and is a subspecies of the dynamic version. If, on the other hand, participants are allowed to 'slide' back and forth across the continuum (again one at a time and in serial order (except at the "jumping" points)), we do not have a forced-march version, but still a dynamic version of the paradox.

[14] See Soames (1999).

[15] This is Terrence Horgan's term. See Horgan (1998) for details.

to the ordinary version of the paradox (CS): There is no single context in which every conditional in IP is judged true. We tend to think that IP *is* true simply because for every conditional composing IP we examine, it does turn out true. However, given the subtleties of context-shifting, we easily miss that context changes for each individual judgment.

We can demonstrate contextual effects independently of any particular version of the paradox. Suppose we have a deck of one-hundred cards ranging from red to orange. Now, pre-theoretically, everyone will agree to the following intuitively obvious principles:

(i) For any color card in the deck, it is red (orange) if, and only if, it is red (orange).

(ii) For any two color cards in the deck, if the first is more red (orange) than the other one, then if the first is orange (red), then the other is too.

Suppose we give a participant the deck and allow her to order the cards from the card she considers to be the most red to the card she considers to be most orange. After documenting her initial ordering, suppose we take this deck and replace the twenty-first card with a card she would judge to be borderline red, say her original forty-fifth card. We then present the first twenty cards one at a time. All will be clearly red. Now we present her with the twenty-first card. Given that she has judged the first twenty cards red, and given that there is a relatively significant difference not only in color between the twenty and twenty-first cards but also in the transition in the degree of redness, our subject may very well judge the twenty-first card to be orange. (The same result would follow if we begin with the last card and work backwards, having the twenty-first card be a borderline case she originally judged orange.) Now suppose we continue just as before, leading her through the series she had previously ordered. Eventually she will reach her original forty-sixth card and judge it to be red. As such, our subject has violated the intuitively obviously principles she pre-theoretically accepts.[16]

We cannot explain these phenomena without contextual considerations. If we revealed to our subject that she has violated these pre-theoretically obvious principles, she might retract the judgments violating those principles. However, she will feel somewhat slighted in doing so. After all, it was loyalty to immediately prior conversational commitments that prompted her judgments. She might (rightly) complain that, in tricking her, we were taking advantage of ordinary, cooperative communicative practice. That is, by tolerance, she cannot judge adjacent pairs differently – she would be violating obvious cooperative conversational principles if she did, but more on this shortly. Since the first twenty cards we presented were adjacent in the series, she has no choice but to classify them alike. But after twenty cards saliently similar (pairwise) in their difference in degree of red, a certain standard for judging upcoming cards has been set. She naturally expects the twenty-first card to fall under this same standard. Since it does not fall under this standard, meaning there is a significant difference in the degree of red between the twenty-first and twentieth in comparison to all cards preceding it, she naturally tends to classify this (already) borderline red card differently than the ones preceding it. But then her judgment concerning the twenty-first card has everything to do with which cards preceded it, and not where it fell in the original spectrum. In short, we tricked her by establishing a certain (conversational) standard and then exploiting it. And tricking others by exploiting operative conversational standards is not ordinarily considered cooperative behavior.

Clear Cases and Borderline Cases

---

[16] I sincerely thank Hud Hudson and Yuri Balashov for a related discussion.

We can divide predications involving clear cases and predications involving borderline cases into further sub-cases based on familiarity or unfamiliarity with the object in question. In the first case, all conversational participants are familiar with the object under discussion. This happens when, for example, the relevant object is in plain sight, is a mutual acquaintance, or is something assumed to be recognized by everyone (e.g. celebrities, important historical figures, natural wonders, *etc*). In the other case, it is not assumed that the object under discussion is familiar to all conversationalists. We will examine the former case first.

In general, for any clear (non)case of a predicate, we feel confident that it (does not) satisfies the predicate. We have no doubt that Kojak is bald and Danny Devito is not tall. However, when the relevant object is familiar, we do not normally make assertions using clear (non)cases except for corrective or didactic purposes. For instance, if we are looking at a man with no hair on his head, it would be infelicitous for me to assert that he is bald. I would do so only to correct your mistaken judgment that he is not bald, or perhaps to teach you the meaning of ‘bald’. Such infelicities result from violating Grice’s Maxim of Quantity – make your contribution as informative as is required.[17] Generally speaking, we do not assert something unless we think it will be informative, and what is obvious is not informative. For this reason, assertions involving familiar clear (non)cases not intended for corrective or didactic purposes typically signal conversational implicatures. That is, we immediately search for alternative possible meanings in order to make the statement non-trivially true.

On the other hand, it is in some significant respect constitutive of the meaning of ‘borderline case’ that members of the linguistic community could disagree on whether the thing in question satisfies the predicate. It is for this reason that we find it dogmatic to insist that a familiar borderline case definitely (does not) satisfies the predicate. As Soames has emphasized,

[17] Grice (1989, pp. 22-41).

there is a kind of room for maneuvering within borderline cases of vague predicates.[18] Here, to be a borderline case just means that we could legitimately judge the object as one way or the other. The sense of dogmatism, then, derives from neglecting this kind of linguistic freedom. *Borderline-borderline cases* (i.e. objects for which it is unclear whether they are clear (non)cases of *F* or borderline cases of *F*) present a problem here, though. In such cases it is possible for one person to legitimately maintain that the object clearly satisfies the predicate while someone else legitimately denies this.

Given Grice's maxim, it follows that for familiar objects the majority of predications involve borderline cases. Furthermore, since borderline cases are by their nature debatable, this means that a speaker making such an assertion does not take her audience to automatically agree with her judgment. In other words, the speaker does not expect her audience to take the truth of her assertion for granted. Rather, she is prepared for her audience to disagree with her judgment. This explains why these predications are typically qualified by hedges such as "I think that…" or "It seems to me that…".

When the object is not assumed to be familiar, we can use assertions involving clear (non)cases informatively. I might describe John to Ann (who is unfamiliar with John) as tall, thin, and bald. Ann, assuming I am being sincere, must also assume that John is a clear case of these three vague predicates (at least to me). For if he were not a clear case, then it would be debatable as to whether John satisfied the relevant predicates. But then I would have violated Grice's Maxim of Quality – try to make your contribution one that is true. Hence, it is clear that

[18] Soames says the following in regard to borderline cases: These objects "are those about whom the semantic rules of the language governing the predicate issue no verdict. However, this does not mean that the predicate can never correctly be used to characterize them. Rather, this is a realm of discretion reserved for individual speakers and hearers. If on a particular occasion one wishes to characterize an individual x in the intermediate range of bald, one is free to do so provided that others in the conversation are prepared to accept this characterization." (1999, pp. 210)

my description is informative, and it is informative in virtue of John being a clear case of the relevant predicates (at least to me).

On the other hand, we do not make normally make assertions involving borderline cases of unfamiliar objects. Describing John as borderline tall would be less informative than describing him as definitely (not) tall. This evidenced by the fact that even though John's being definitely not tall (semantically) entails that he might be a borderline case, if I were to assert this, Ann would automatically assume that John is probably short. The best explanation here is that we try to interpret others so as to make what they say true, and we do not take assertions involving borderline cases as unconditionally true.

The obvious difference between the two general cases lies in what the conversational participants believe and what they believe others believe. When the object is familiar, speakers will have certain beliefs about that object and they will assume that other speakers share some of those beliefs as well. If you and I are looking at a man with no head hair, it would be redundant for me to assert that he is bald since I believe he is bald and believe that you believe this as well. On the other hand, if we are looking at a genuine borderline case of bald, it would be dogmatic for me to *insist* that he is bald since, even though I might think he is closer to bald than not, I have no reason to think that you should believe this as well. When the object is unfamiliar to the hearer but not the speaker, the speaker naturally assumes that the hearer has no relevant beliefs about that object. In the case described, I believe that Ann has no beliefs concerning John's height, shape, or hair situation. As such, I can informatively, unqualifiedly assert that John is tall, thin, and bald only because he is a clear case of all three predicates (at least to me) and, since Ann is unfamiliar with John, I have reason to believe this is not obvious to Ann. However, if John is a borderline case of bald, it would be odd for me to describe him as such since it is constitutive of ‘borderline case’ that we could disagree on whether John is bald, and Ann is in no position to have any such relevant beliefs.

Going along with Robert Stalnaker, let’s call a proposition *p* a *pragmatic presupposition* if (i) all conversational participants in a conversational context *c* believe that *p*, and (ii) all believe that everyone else in *c* believes that *p* as well. Let’s call the set of all these presuppositions along with all assumptions, ordinary presuppositions, *etc* (i.e. all propositional attitudes operative in a conversational context) accepted by all conversational participants *common ground*.[19] In terms of common ground, assertions involving familiar clear (non)cases are redundant simply because they are common ground presuppositions. Oppositely, assertions involving familiar borderline cases are never common ground presuppositions, and this is why insisting that such an object really does (not) satisfy the predicate is dogmatic. Likewise, by definition, predications involving unfamiliar objects are never common ground presuppositions.

The crucial difference, however, is that we take predications involving unfamiliar clear (non)cases as informative, i.e. we learn something about the object, whereas we do not take predications involving familiar borderline cases to be informative. This difference is exhibited in the fact that once Ann believes that John is a clear case of bald, she will carry this belief over into subsequent conversations. Contrarily, it would be absurd to think that agreements concerning familiar borderline cases hold outside of that particular context. For example, suppose A and B are looking at borderline bald C and D, where D has slightly less hair than C. Suppose A and B agree to count C as bald. In doing so, they have thereby committed themselves to counting D as bald as well.[20] However, tomorrow, when A is with E instead of B, it would be absurd for A to judge D bald based on his prior agreement with B concerning C. Clearly, we

[19] For discussions relevant to pragmatic presuppositions and common ground, see Stalnaker (1974) and (1999).
[20] This is an example of what supervaluationists call *penumbral connections*, or certain obvious inferences based on certain background assumptions. See Fine (1975) for an explication.

recognize that these kinds of judgments hold good only within the conversational contexts in which they are established. This means that we treat judgments about familiar borderline cases as temporary assumptions operative only relative to a context.

Denying the Inductive Premise

Most theorists of vagueness are committed to denying IP. I would like to understand how we can deny IP without being committed to unconditionally denying obvious truths like the following:

(BIP) A single hair doesn't make the difference between the bald and non-bald men.

(TIP) There's no sharp boundary between the tall and non-tall people.

(HIP) We can't go from a heap to a non-heap by taking away a single grain of salt.

These all say the same thing, namely that the respective inductive premises for 'bald', 'tall', and 'heap' are true. Incredibly, I will argue that sometimes we can legitimately deny all three of these without denying their obvious plausibility. Moreover, I will argue that we never unconditionally take their truth for granted.

It is platitudinous that conditional truths do not sum up to unconditional truths. If I assume that *p*, *q*, and *r* and thereby deduce *s*, I cannot thereby legitimately conclude *s*. Rather, I can only legitimately conclude that *if p*, *q*, and *r*, *then s*. Likewise, if I judge some range of borderline cases $\langle o_i,\ldots,o_j\rangle$ as bald, I cannot thereby conclude unconditionally that $Bo_i,\ldots,Bo_j$. Rather, I would conclude that $Bo_i,\ldots,Bo_j$ only relative to some context *c*. As indicated, it would be absurd for me to assume these judgments hold for all contexts. Furthermore, the truth-status of a universal generalization like IP is collectively determined by the truth-statuses of its component conditionals.[21] These platitudes clearly indicate that the inductive premise of any legitimate Sorites, i.e. one which meets S1–S4, is not unconditionally true. And this seems to fly in the face of obvious truths like BIP, TIP, and HIP.

Now, not all inductive premises are equal. Some are more plausible than others, some being outright false, others very plausible, others still we might call borderline-plausible. Consider Sorites' for 'tall' partitioned by meters, inches, and micrometers. For all three of these Sorites, we have the same inductive premise, namely

(TIPS) $(\forall x)(To_i \rightarrow To_{i+1})$

In all three cases, the inductive premise reads the same, namely that there is no sharp boundary between the tall and non-tall people (TIP). However, the "Sorites" partitioned by meters fails to be paradoxical at all – its inductive premise is plainly false. Contrarily, the Sorites partitioned by micrometers is very plausible – surely we cannot deny it. Finally, the Sorites partitioned by inches is a borderline case. Some might it find plausible while others might not. For some reason or another, someone might have it fixed in their mind that any American male 6'4" in height or more is tall. For them, this IP is plainly false since there will be a sharp boundary. Someone with a more liberal conception of tallness might find this inductive premise somewhat plausible though.

Furthermore, individual inductive premises undergo contextual effects just as vague predications do. Suppose we have someone evaluate twenty IP's partitioned by micrometers and measurements close to micrometers. Now suppose we present them with the IP partitioned by inches. Even if they normally would be disposed to judge this one as plausible, they may now find it unpersuasive just as the person presented with twenty clearly red cards might be prone to

---

[21] I take exception with supervaluationism here. Supervaluationists maintain that an existential generalization can be true even if none of its instances are.

call the borderline red card orange. What I am describing is just *higher-order vagueness*, vagueness as to what is to count as vagueness. It shows that inductive premises, like vague predications, are both subject-relative and context-dependent. Moreover, it shows that we do not take the truth of inductive premises like TIPS for granted. We must know something more about them before we judge whether or not they are true. In particular, we must know how finely-partitioned the domain is before we are in a position to evaluate the inductive premise.

What determines the plausibility of these three IP's sharing the same English reading, TIP? Again, nothing about TIPS itself determines the plausibility of TIP – it reads as TIP in all three cases. TIPS partitioned by meters fails to be paradoxical simply because it divides the tall and non-tall people into clearly distinguishable sets. Contrarily, TIPS partitioned by micrometers divides the domain into so many sets that we cannot legitimately claim that a sharp boundary exists between one but not another. In the intermediary case, the partitioning might be fine enough to be plausible for some people but not others. Out of the three cases, only micrometers satisfy S3. Meters obviously do not satisfy S3, and inches are a borderline case. Since S3 (there is no discernible difference between any object-pair) ensures that object-pairs satisfy WP (indiscernible pairs must be classified alike), this strongly suggests that WP is actually responsible for the plausibility of inductive premises.

My claim is that when we read the plain-English BIP, TIP, and HIP, we implicitly assume that they already meet WP and thereby grant their obviousness. The point is easier to make with 'tall' because we can choose any way of measuring we like. It is less obvious with 'bald' or 'heap', but if we use our imaginations we can find falsifying instances of them too. We are used to head hairs being of a similar diameter, one so small that no single hair could possibly make the difference between being bald or not. However, if we think of head hairs of abnormally large diameters, say one hair that covers the entire top of the skull and three which cover the sides and back of the crown, then we may very well find BIP false. Likewise for HIP if we think of abnormally large grains. I admit that these are strange cases, but I see no reason to accept that a necessary feature of the meanings of 'head hair' or 'grain' is their being of a certain limited diameter and size. Even if there is a reason for accepting this, surely it is vague as to what diameter or size this is anyway.[22]

The fact that these are difficult scenarios to imagine further bolsters my claim that we tend to presuppose BIP, TIP, and HIP already satisfy S3 (and consequently WP). Upon reflection, this is precisely why 'heap', 'tall', and 'bald' serve as prototypical Sorites predicates. As for 'tall', we typically do not measure human height by whole meters. Doing so would not be very informative, for it would not distinguish individuals very well. Similarly, we do not typically measure height by whole micrometers because doing so would be impractical and too informative – how does one who measures fifty micrometers taller than another compare to the average person's height? However, it is easier to imagine measuring height by different intervals than measuring numbers of hairs of abnormal diameters or grains of abnormal sizes. This is why 'tall' is a better Soritical predicate for our purposes. Importantly, we should recognize that assuming that BIP, TIP, and HIP already satisfy S3 is an easy mistake to make. After all, it would be rather misleading if, in explaining HIP to someone, I was really assuming anything like abnormally large grains of salt.

[22] The situation is perhaps even more difficult with color predicates since they naturally divide into a fluid spectrum. If, instead of thinking of the relevant domain of 'objects' as points on the spectrum, we think the transition from (say) red to orange by color cards, the analogy goes through easily enough. The more cards, the more finely-grained the partitioning, and *vice versa*. For example, if there were only four cards where there is an equal difference in redness between adjacent pairs, then intuitively we have no paradox since the partitioning would fail to satisfy S3.

Why the Inductive Premise is Never Unconditionally True

I said that *no* inductive premise is unconditionally true, even the plausible ones. This seemed to fly in the face of obvious truths like BIP, TIP, and HIP. The plausibility of these obvious truths was then to be grounded in their alleged compliance with S3. And, recall, S3 guarantees that the relevant predicate satisfies WP. I now claim that WP is a common ground presupposition, i.e. a truth taken for granted in any conversational setting and believed by all to be so. Jamie Tappenden has noted that we cannot better explain the attractiveness of a legitimate inductive premise to someone who does not understand its attractiveness than by stating something like the following: "Look, if two objects are (for the purposes of ordinary conversation) observationally indistinguishable, then one is red if, and only if, the other is too."[23] Someone sincerely and consistently denying this claim would rightly be accused of linguistic incompetence. They do not understand the meaning of 'red', or they do not understand the meaning of the biconditional.

One persuasive argument here is that we learn the meanings of observational predicates by ostension. Generally, ostensive predicates are learned via prototypical instances. We then extrapolate from those initial ostensive instances in applying the predicate to further instances. So judging two ostensively identical objects differently belies how one actually learns the predicate in the first place. As Tappenden has argued, we typically would utter something like WP only for corrective purposes. In doing so, we are attempting to realign the mistaken person's judgments with that of the linguistic community's. Put differently, we are trying to bring the

[23] Tappenden says the following: "If someone consistently demonstrated a misunderstanding of the proper use of 'looks red to me' or 'observationally indistinguishable', one might well say to him: "Look, if any two samples are observationally indistinguishable to you then one looks red to you if and only if the other looks red to you."…[T]he sentence will probably have the desired effect of getting the hearer to stop talking in the proscribed way." (1993, pp. 572-3)

mistaken person's linguistic use into conformity with linguistic convention. And this, in turn, makes sense only within the framework of how we believe all users of the language would consistently use the predicate. For suppose Ted sincerely and consistently judged a clear case of green to be non-green. We could give him a color chart, but in doing so we would be implicitly appealing to linguistic convention since it is how we (*viz*. the linguistic community as a whole) actually use the predicate which determines which color patch is associated with 'green' on the chart. In short, in appealing to linguistic convention we are appealing to common ground, the presuppositions we believe to be presupposed by all. The fact that we would utter something like WP only for corrective purposes shows that we presuppose its truth and presuppose that others presuppose it as well.

Since WP is common ground, it is little wonder that we find BIP, TIP, and HIP attractive. Since we are apt to assume that they already satisfy S3, and S3 (along with the other constraints) guarantees that the predicate satisfies WP, we naturally grant their obviousness. But what about the individual conditionals composing IP? We should consider this question by first considering the truth-status of the predications constituting their antecedents and consequents. For any genuine IP, there will be clear (non)cases and borderline cases. Again, predications involving clear (non)cases will constitute common grounds presuppositions. This alone guarantees that the objects satisfy S4. On the other hand, predications involving borderline cases by their very nature cannot be commonly presupposed. We grant their truth (falsity) only on a conditional basis and relative to a context.

Now, by S4, we will have no choice but to grant the truth to the first predication and a "suitable range" following it. Likewise for the truth of the negation of the last predication and a "suitable range" preceding it. Notice, however, that this holds independently of any conditional we might form using those predications. For example, for any borderline case o_i, the following conditionals will be (classically) true[24]:

(4) $(\neg)Fo_i \rightarrow Fo_1$

(5) $(\neg)Fo_i \rightarrow \neg Fo_n$

(6) $\neg Fo_1 \rightarrow (\neg)Fo_i$

(7) $Fo_n \rightarrow (\neg)Fo_i$

Obviously, (4) and (5) are (classically) true in virtue of their consequent's truth, and (6) and (7) because of their antecedent's falsity. Moreover, for all clear cases $<o_1,\ldots o_i>$ and clear non-cases $<o_j,\ldots o_n>$, we would take (8) and (9) to be true apart from any considerations about conditionals they might form:

(8) $Fo_1 \wedge Fo_2 \wedge ,\ldots, \wedge Fo_i$

(9) $\neg Fo_j \wedge \neg Fo_{j+1} \wedge ,\ldots, \wedge \neg Fo_n$

Likewise, we would take any conditional formed out of predications exclusively in (8) or (9) to be (classically) true even if the predications composing those conditionals were not adjacent in the series. In short, the truth of each conditional is entirely owing to the truth of its components. Consequently, the truth of any of these conditionals is in no way owing to its antecedent and consequent being adjacent members in the series.

Not so with conditionals composed out of borderline cases. Given that we take borderline cases as conditionally true only, we cannot judge the truth of (10), where $o_i,\ldots,o_j$ are borderline cases, independently of a particular context:

(10) $(\neg)Fo_i \wedge (\neg)Fo_{i+1} \wedge ,\ldots, \wedge (\neg)Fo_j$

24 I adopt the convention that for any predicate P and variable v, '$(\neg)Pv$' means Pv could be either true or false.

However, of course, we can and *do* judge (11) true in any particular context (for borderline object o_i and genuine Sorites predicate F):

(11) $Fo_i \leftrightarrow Fo_{i+1}$

What allows us to do this? By now, the answer should be obvious. By S3, (11) satisfies WP, and since WP is presupposed to be true, so does (11). Hence, we see that the truth of any single conditional composed of predications of adjacent borderline cases is entirely owing to the obviousness of WP, *not* the truth of its components. This means that once we recognize that some IP satisfies S1–S4, we need not even check to see if any of the composite conditionals are in fact true. Of course, they will be true if we do check them, but we can know this independently of how we might judge them in any particular conversational context.

We can further convince ourselves of WP's culpability if we revisit the Sorites for 'tall'. In the micrometer case, since the domain is so finely-partitioned, nothing about adjacency in the series contributes to the truth of a given conditional. Clearly, we will find (11) true in this case, but we would also find it true if it were (11') instead:

(11') $Fo_i \leftrightarrow Fo_{i+1000}$

If Sorites can be legitimately constructed out of millimeters, then surely one constructed by one-thousand micrometer increments is also legitimate. Equally obvious is the fact that the same cannot be said for a Sorites partitioned by millimeters. One-thousand millimeter increments are surely large enough to legitimately draw a sharp boundary. The only explanation here is that former satisfies WP while the latter does not.

But this means that we must admit that the inductive premise of a legitimate Sorites *is* in fact an unconditional truth after all, right? No. Again, the truth-status of a universal generalization depends on the truth-status of all its instances. In actual practice, we treat predications of borderline cases only as temporary assumptions. As such, they are only conditionally true (false) relative to a context. Likewise for conditionals composed of those predications. Since IP quantifies over *all* relevant conditionals, it cannot be unconditionally true, a common ground presupposition. Rather, as we have seen, its plausibility is owing to WP. As a result, we do not take the truth of any inductive premise for granted. This is evidenced by the fact that in order to grant its truth, we first need to know whether it satisfies S1–S4. But this along with WP alone is enough to guarantee its truth independently of evaluating any of its component conditionals. We easily neglect this fact by dealing only with prototypically vague predicates, ones like ‘bald’ and ‘heap’ for which it is difficult to imagine abnormal interpretations of those predicates along the relevant dimension, e.g. understanding by ‘head hair’ anything like ‘a hair covering the entire top of the skull’.

Section 4: Efficiency and the Sorites

The Semantics of Scalar Adjectives

When introducing the Sorites to someone unfamiliar with the paradox, a common response is something like "I just don't understand how the paradox is relevant to the meaning of 'bald'. After all, by 'bald' we mean more than just number of hairs." This section is an attempt at codifying this commonsensical intuition. I will argue that a significant problem with the Sorites is singling out a particular semantic dimension of multidimensional prototypical vague predicates like 'bald' and 'heap'. Furthermore, given arguments from the last section, even with seemingly unidimensional vague predicates like 'tall' and 'expensive', there is no straightforward transition from the (generally accepted) semantics of these predicates to legitimate Sorites series. The upshot is that the Sorites reveals no obvious problems with the semantics of our ordinary vague predicates. I will argue that for evolutionary game-theoretic considerations made earlier, we should not expect the semantics of our everyday predicates to be as fine-grained as is required for legitimate Sorites construction. Rather, our interest in the Sorites must ultimately be set-theoretic in nature, i.e. our fascination with the Sorites is really owing to the problems it poses to classical logic and semantics.

I ended the last section by claiming that we tend to ignore that not all Sorites are legitimate by focusing on certain prototypical vague predicates for which non-standard interpretations are difficult. This is not surprising when we consider what kind of predicates we have been dealing with, namely *scalar adjectives* (or *gradable adjectives*). Nothing is tall simpliciter, but rather in comparison to some other relevant class of objects. To call something "tall" is to locate it somewhere on a scale of height in comparison to other things. This set of compared objects is commonly called a *comparison class*. Further, the comparison class for 'tall' and other scalar adjectives commonly fluctuates. Depending on context, (12) could mean a number of things, e.g. (13), (14), or (15):

(12) Jim is tall.

(13) Jim is of greater than average height for an adult American male.

(14) Jim is of greater than average height for a professional basketball player.

(15) Jim is of greater than average height for a midget.

The intuition is that (16), (17), and (18) can be true, meaningful sentences which cannot be explained via some kind of objective, absolute context-independent properties:

(16) Eric is of average height for an adult American male, but he was tall for a Taiwanese.

(17) Ted is bald, but not in comparison with Kojak.

(18) Both pile A and pile B are heaps of salt, but A is more of a heap than B.

If we suppose there are objective properties "tallness", "baldness", or "heapness", it seems that in order to make sense of (16) – (18) we need to say that objects have them and do not have them depending on context. For instance, I am not now an instance of tallness, but while living in Taiwan I was. Similarly, if (17) is not to be straightforwardly contradictory, it seems that we must say that Ted is an instance of baldness when considered in isolation, but not in comparison with balder objects. Finally, it is not entirely clear how a non-comparative based semantics of scalar adjectives would make sense of the second conjunct in (18). Does heap A "instantiate baldness" more than heap B, but in a non-comparative way? Given these strange metaphysical commitments, current semanticists of scalar adjectives commonly assume that the arguments of such predicates combine with contextually given comparison classes to deliver a position on a semantic scale (composed of those comparison classes). Viewing prototypical vague predicates this way, it becomes clear why we tend to assume that inductive premises like BIP, TIP, and HIP are legitimate: The contextually determined comparison classes for 'bald' and 'heap' are never actually constituted in the way imagined, i.e. by four hairs covering the entire skull or abnormally large grains of salt. As for 'tall', for obvious practical reasons, we do not divide human height by meters for comparative purposes.

To elucidate this common assumption operative within the semantics of scalar adjectives, we will need to draw the broadly Kaplanian distinction between the *character* and *content* of expressions.[25] In general, the character of an expression is the linguistic rule determining its semantic value, or its content. The distinction is perhaps best illustrated via indexicals. For example, the character of 'I' is something like 'the speaker of utterance'. According to Kaplan, the character of an expression combines with context of utterance to deliver its semantic content. This explains how different utterances of (19) could express different propositional contents, e.g. (20) or (21):

(19) I am here.

(20) Eric Snyder is in Espresso Royale coffee shop.

(21) Ted is in the living room.

Applying Kaplan's distinction to the vague predicates of interest, the following seem like good candidates for their respective characters:

(22) the character of *tall* = having significantly greater than average height relative to some comparison class *C*

(23) the character of *bald* = having significantly less head hair arranged in a certain manner relative to some comparison class *C*

---

[25] See Kaplan's classic (1989) for the character / content distinction. The semantics of scalar adjectives and gradable predicates in general has occupied semanticists since (at least) the early 1970's. For a nice representative sample, and one that has direct bearing on what follows, see Kennedy (1999).

(24) the character of *heap* = being a significantly larger than average collection of some granular substance arranged in cone-like shape relative to some comparison class *C*

Now, as with indexicals, the character of an expression combines with context to deliver semantic content. However, instead of delivering some directly referring singular term as in the case of indexicals, for scalar adjectives context provides the relevant comparison class, as required for making sense of the variability of sentences like (12). Linguists further impose a total linear ordering on the contextually given comparative class, thereby delivering the requisite semantic scale. Importantly, this ordering is based on the degree of the general property in question. The character of the expression then locates the argument's position on this scale. The addition of "significantly" in (22) – (24) is required to make sense of sentences like (25):

(25) Jim is of greater than average height for a midget, but he is not a tall midget.

If being 'tall' were merely being located above the median value on the semantic scale, (25) would be contradictory. But, intuitively, there are contexts in which (25) could be used felicitously. Another reason for adopting the "significantly" clause comes from what Dianna Raffman has termed *crisp judgments*.[26] Suppose we have two philosophy books, A which is 501 pages long and B which is 500 pages long. Now consider (26) – (28):

(26) A and B are long.

(27) A is longer than B.

(28)* A is long compared to B.

Since A and B are very similar in length, both are either long or not long. Since philosophy books typically tend not to be so long, and assuming that the relevant comparison class here is philosophy books, (26) is true. This also explains the truth of (27) – no matter how large or

[26] See Raffman (1994).

divergent the semantic scale happens to be, since A is located above B, the sentence is true. But it is difficult to imagine contexts in which (28) could be felicitously uttered. The best explanation seems to be that since the respective lengths are so similar, there are no relevant contexts for which a difference of a page could be significant.[27]

This intuitive account of scalar adjectives easily accommodates the previously problematic (16) – (18). In (16), I am located toward the middle of the semantic scale of height composed of Americans, whereas I am located significantly high on the scale composed of Taiwanese. (17) first locates Ted and Kojak at the lower end of the baldness scale, and then locates Kojak significantly lower than Ted; if we then take the lower end of the original scale as a new comparison class *C'*, Ted will be toward the upper end of this new scale and Kojak toward the bottom. Finally, (18) locates both A and B toward the upper end of salt pile sizes, and then places A above B on that scale. Notice that in all cases, we must appeal to context-shifts to account for the shifts in comparison classes. That is, since new comparison classes are required to make sense of felicity of (16) – (18), and since we must suppose that context delivers comparison classes to explain sentences like (12) can express sentences like (13), (14), or (15), we must appeal to context-shifting in the sentences involved. The upshot of the proposal is that scalar adjectives semantically resemble indexicals more closely than one might initially think – their semantic content (i.e. their place on the contextually determined semantic scale) depends entirely on features of context.[28]

---

[27] It is important to recognize that the account given for scalar adjectives is general enough to incorporate multiple popular theories of vagueness based on comparison classes. The two most obvious contextualist candidates are Graff Fara's (2000) and Raffman's (1994). The account is also general enough to incorporate a supervaluationist version of context-dependent comparison classes. The best candidate is Kamp and Partee's (1995) prototype theory. Here, we begin with contextually determined conceptual prototypes, which serve as polar values for comparison classes, and perform precisifications based on these.

[28] I intentionally avoid the obviously problematic nature of semantic scales. For example, on Kennedy's treatment, there is a definite point dividing the tall from the non-tall. By analogy to higher-order vagueness, it would seem that this supposition is *prima facie* unwarranted. I tend to agree with Heck's (2003, pp. 111-13) comments about scalar

Why the Sorites Reveals Nothing of Semantic Interest

How is any of this relevant to the Sorites? Most importantly, notice that the structure of the contextually determined semantic scale very closely resembles that of Sorites continuums, i.e. the ordering of objects according to S1 – S4. It is little surprise then that scalar adjectives serve as prototypical vague predicates. We should *not*, however, mistakenly identify semantic scales with Sorites continuums. Pointing out some significant differences between the two should suffice. First, contrary to S1, the semantic scales of scalar adjectives need not be based on some single dimension. Consider 'bald' for instance. At least two semantic dimensions, number of hairs and arrangement of hairs, are relevant to determining the semantic scale. Suppose Fred, who has *n* hairs, is bald. Now gradually transplant these *n* hairs (in an even distribution) to the top of Fred's skull. Depending on *n*, we would be less likely to call Fred 'bald'. Also, depending on how we distribute the hairs, there might be independently constructible Sorites'. Similarly, consider a heap of salt with *n* grains piled on a large flat surface. Now gradually flatten the pile, spreading all *n* grains evenly across the surface. We would not normally call this a 'heap'. Again, however, there will be possible Sorites' constructible from the gradualness of the flattening. The multiple dimensions of scalar adjectives must collectively determine their respective semantic scales, and while Sorites can often be constructed out of individual dimensions, there is no reason to suppose the ordering of semantic scales via multiple dimensions mirrors the orderings

---

adjectives: "Such adjectives are associated with scales…and the context-dependence of such adjectives is explained in terms of the fact that 'tall means roughly: of a degree of height greater than δ, where δ is a contextually determined degree of height. Context is thus obliged, in any case, to fix a point along the scale that will divide the tall from the not-tall…I simply see no reason to suppose that ordinary contexts fix unique such degrees, nor even that they fix the degrees precisely to decide, of every object in some contextually relevant domain, whether it counts as tall or not. What I suspect, rather, is that context *restricts* the set of degrees as far as is needed for conversational purposes and that further such restrictions are negotiated as they become necessary…If context is insufficient to decide whether, say, Bob counts as tall in it, if it *matters* whether Bob counts as tall, there will be a problem. But it need not matter, and there need be no problem."

of individual dimensions. Indeed, we should expect this to not be the case – hence, the counterexamples just considered.

Secondly, contrary to S3, semantic scales need not satisfy tolerance. Consider a scale formed by three individuals, the first 4'6", the second 5', and the third 5'6". This comparison class clearly fails tolerance but easily forms a scale by imposing the proper ordering. Finally, contrary to S4, semantic scales need not contain clear (non)cases at the poles. Suppose we are looking at ten borderline-tall men (for adult American males) and we are interested in determining who is tall relative to this group. We can order the men according to their height, thereby determining a semantic scale and relative tallness from it. However, none of the members constituting the scale are clear (non)cases of 'tall' (relative to adult American males), and if the members of the new scale vary only minimally in height, there may not even be clear (non)cases of 'tall' relative to this subgroup either.

If the Sorites is to reveal anything semantically interesting about our prototypical vague predicates, we should expect that either (i) Sorites can be formed directly out of the relevant semantic scales, or (ii) semantic scales convert smoothly to Sorites continuums. However, I submit, neither of these expectations is generally fulfilled. The first problem results from the *multidimensionality* of many vague predicates. We cannot legitimately construct Sorites' from number of hairs *along with* distribution of hairs simpliciter. Instead, we must choose some single dimension and a unit for partitioning that dimension. But, in general, different semantic dimensions will require different kinds of partitionings and different units for those partitionings – witness 'bald' and 'heap'. Likewise, we cannot construct legitimate Sorites continuums from height simpliciter. That is, even (seemingly) unidimensional vague predicates like 'tall', 'expensive', or 'old' do not transition immediately to Sorites continuums. This is because we must first choose a unit for partitioning – micrometers vs. feet, cents vs. thousands of dollars, seconds vs. years. And, as emphasized in the second part, which unit we choose is not arbitrary. In fact, the only way to preserve the semantic scale is to choose the smallest unit in which every element in the semantic scale receives its own cell in the Sorites partitioning.[29] And given a sufficiently large domain, the chosen unit may need to be very small indeed – consider, for instance, partitioning the domain of adult American males by height. The worry then, of course, is that there is no general guarantee that all possible cells in the partitioning will receive a representative. In such a case, there is no guarantee that IP will be false, since falsifying IP requires a boundary-establishing pair $<o_i, o_{i+1}>$, and if there are gaps in cell representation, it may turn out that either o_i or o_{i+1} is empty. Similar problems plague the second expectation. For multidimensional predicates, there is simply no guarantee that the ordering in the semantic scale mirrors that of a (unidimensional) Sorites continuum. So we cannot in general abstract Sorites continuums from semantic scales; some kind of reordering will be required. I conclude that, assuming something like the contemporary semantics for gradable adjectives is correct, the Sorites reveals nothing obviously problematic with the semantics of prototypical vague predicates. As a corollary, contrary to nihilists like Wright and Horgan, the Sorites does *not* reveal incoherence in natural language predicates.[30]

However, even if there were a way of smoothly importing semantic scales to Sorites continuums, there is an important sense in which the resulting Sorites would be of little interest. Let's consider how this would happen. First and foremost, according to S1, we must isolate a single semantic dimension. As mentioned, the ordering of a semantic scale need not mirror the ordering of the dimension we choose. So we will need some way of reordering the objects

---

[29] Barring ties, of course. For example, A and B have, down to the very last penny, identical wealth.
[30] See Wright (1975) and Horgan (1998).

composing the comparison class to be in accordance with S1 and S2. Let us (dubiously) assume there is a function *f* from *C* to $<o_1,\ldots,o_n>$ of the conditional Sorites which does just this. Importantly, recognize that once this has happened, and assuming there is a large enough domain of objects and at least some difference in degree of *F*-ness between them (along the relevant dimension), we will have in some sense trivially satisfied S4. It will be trivially satisfied by the comparative nature of the scalar adjective we are interested in. That is, given the total linear ordering imposed by S1 and S2, there will trivially be a greatest member and least member of the new series. Since for the purposes of the Sorites we are in a sense defining 'tall', 'bald', 'heap', *etc* by the ordering imposed on the dimension chosen, there will trivially be a tallest / least tall, baldest / least bald, most heap-like / least heap-like, *etc* in the series. But since objects are tall, bald, heap-like, *etc* only relative to some comparison class of objects, this means that we can take o_1 and its immediate successors as *F* relative to *C* and o_n and its immediate predecessors as non-*F* relative to *C*. In this way, S4 is trivially satisfied. Of course, this is not what we normally have in mind when constructing the Sorites. Normally, we think of o_1 and its immediate successors as *objectively F* and o_n and its immediate predecessors as *objectively* not-*F*. But this is to ignore the essentially comparative character of scalar adjectives surely necessary for adequately accounting for the semantics of sentences like (12) – (18). Upon recognizing this, we must admit that (say) a Sorites for 'tall' constructed relative to the domain of midgets is as equally legitimate as the one (say) constructed relative to the domain of adult American males. If this result strikes you as fundamentally misguided, since perhaps you think midgets are definitionally not tall, ask yourself: Not tall relative to what? *There are no objective, non-context-sensitive properties of tallness, baldness, etc*.[31]

[31] This might remind the reader of Unger's (1979) nihilism regarding the inexistence of objectively flat objects.

Graff Fara's Interest-Relative Account

So far, this account accords nicely with a particular brand of contextualism, namely Delia Graff Fara's (2000). According to the latter, what counts as 'significant' in a given context depends on the interests and purposes of the conversational participants involved. In terms of the framework given above, this means that where an object gets located on the semantic scale is a function of the comparison class picked out in a context and the interests and purposes of conversationalists. We might then think of the contextually determined comparison class as setting a certain conversational standard for *F*-ness, and interests and purposes as locating an object on the semantic scale determined by that standard. Furthermore, she adopts a tolerance-like principle for contextual significance based on operative conversational standards, what we she calls *the similarity constraint*:

(SC) For any standard *s* operative in the use of a vague predicate *F* in a context *c*, if o_i and o_j are saliently similar according to *s* in *c*, then Fo_i iff Fo_j

This intuitive principle gives a nice explanation for (26) – (28). Since A and B are similar in length, both get classified alike. Moreover, since A and B fall under the similarity constraint in all reasonable contexts, (28) is false. It also affords a nice solution to the Sorites. Graff Fara first keenly notes that variability in conversational standards (i.e. comparison classes) *cannot* alone explain the Sorites. We can fix conversational standards and still find inductive premises plausible. Suppose we fix the comparison class and an average value for this class, say the price of apartments in Athens which range from $200 - $1000 with a $600 average. The following Sorites remains paradoxical despite the fact that the inductive premise should explicitly fail at the pair <$601,$600>:

(29) P1: A $1000 apartment is expensive for an apartment in Athens.

P2: A free apartment is not expensive for an apartment in Athens.

IP: An apartment which is $1 less than an expensive apartment in Athens is expensive

C: A free apartment is expensive for an apartment in Athens.

In other words, even when the standard for comparison is conversationally explicit, we still find IP plausible. So there must be something more to the Sorites than shifting comparative standards. Graff Fara's solution appeals to SC. The very act of searching for a falsifying pair, which happens to be the interest and purpose of the imagined "conversation", renders adjacent objects in a legitimate Sorites series saliently similar. Hence, the very act of attempting to falsify the inductive premise has the odd consequence of being self-defeating. No wonder then that we find the universally quantified IP plausible – every pair we search is classified alike. What is important for our present purposes is how Graff Fara justifies SC. She does so based on cost-benefit considerations. To classify two saliently similar objects differently simply costs more than it benefits, i.e. it is *inefficient*. So even if we *know* that o_i and o_j are in fact qualitatively different, I may still count them as saliently similar for present purposes simply because it is more efficient to do so.

While I believe Graff Fara's theory is the best currently on the market, it suffers from two (potentially crippling) defects, both pointed out by Jason Stanley (2003). First is the *generality problem*. Graff Fara's analysis explicitly concerns the semantics of scalar adjectives, and her solution to the Sorites wholly depends on the semantics of that grammatical category. Recall, speakers' interests and purposes determine the relative position of an object on the semantic scale, itself determined by features of context. Since it is far from clear that other grammatical categories share this semantic feature, Graff Fara's solution threatens to fail to generalize to other Sorites susceptible categories, e.g. 'tadpole', 'adolescent', and other vague nouns. Graff Fara recognizes this worry: "But it is not a semantic feature of nouns that they are associated with a dimension of variation. That is why a generalization of my proposal to nouns would require a case by case analysis."[32] But if vague nouns do not exhibit a similar semantic structure to scalar adjectives, it is difficult to see just how a Sorites constructed from these could be treated by her solution. Stanley's second objection is *the truth-conditions problem*. In essence, Stanley accuses Graff Fara of getting the truth-conditions for propositions containing vague expressions wrong. Since 'significant to *x*' requires an implicit argument, namely the speaker of utterance, and since according to Graff Fara sentences containing vague expressions are essentially interest-relative, (Russellian) propositions containing vague expressions include their subjects as constituents. So, for instance, consider Ted's utterance of (30):

(30) Mount Everest is tall for a mountain.

According to Graff Fara then, (30) expresses the proposition that Mount Everest is for Ted significantly greater in height than the average mountain. But, argues Stanley, (30) would be true even if Ted had never existed and hence had no interests or purposes at all. In fact, (30) would be true if there were no people and hence no interests or purposes at all. So it seems that Graff Fara's implicit relativization of certain propositions to speakers' interests and purposes is too strong.

Whatever the merit of Stanley's second criticism might have, I take the generality problem to be of more pressing concern.[33] In what follows I hope to provide a more general

[32] Graff Fara (2000, pp. 29).

[33] Personally, I am not entirely convinced that (30) would be true even if no humans had ever existed. Stanley's criticism clearly presupposes objective properties of baldness, tallness, *etc*, and as mentioned above this is to ignore the essentially comparative nature of those predicates. If there had been no humans, then presumably there would be

explanation for why we find inductive premises attractive, one which I take to be in keeping with Graff Fara's more fundamental motivation, namely the efficiency of vague expressions. If it is on track, it would account for Graff Fara's particular solution to gradable Sorites predicates as well as other Soritical categories.

A General Solution to the Sorites?

Graff Fara has also given the standard by which theories of Vagueness are to be judged. Any theory of Vagueness must adequately answer three questions concerning the Sorites:

the semantic question: Is IP false?

the epistemological question: If IP is false, which pair establishes the boundary?

the psychological question: If IP is false, why are we so inclined to accept it?

I propose to pay off my explanatory debts via evolutionary game-theoretic considerations. In any actual legitimate Sorites, *if* speakers are forced to establish a boundary, thereby falsifying IP, they can. As Shapiro and others have emphasized, speakers may run back and forth through the continuum, forced by tolerance to classify more and more objects *F* or not-*F*. Eventually, either by exhaustion, indifference, or both, they will pick a boundary-establishing pair. As supervaluationists have long emphasized, the pair chosen will be rather arbitrary. Since, by tolerance, many nearby pairs could have just as legitimately solved the task at hand, speakers will feel somewhat hesitant to choose the pair they choose. If asked why they chose this pair and

---

no comparisons to be made (and obviously no such predicates). As I see it, Stanley's criticism points to a more pressing problem with (Kennedy's) semantics quite generally, and is not unique to Graff Fara's employment of it. *Kennedy's* semantics presupposes objective properties of baldness, tallness, *etc*, and we should certainly question this presumption. But if this is really where the problem lies, then Graff Fara's addition of conversationalists' intentions and purposes is not really the source of Stanley's worry.

not some nearby pair instead, speakers may concede that this other pair may just as well have served as the boundary.

In general, which pair is chosen cannot be determined beforehand, and we should not expect the same pair to be chosen over repeated runs of the same series. What can be given is a general principle relating boundary-choosing behavior and the fineness of partitioning:

(A) The more finely-grained the partitioning of *F*, the more *arbitrary* the boundary-establishing pair chosen will be, and *vice versa*

By "arbitrary" here, I mean unlikelihood of being repeatedly selected on recurring runs of the same Sorites series. Given a sufficiently large domain of men, the likelihood of the same boundary-establishing pair being chosen on repeated runs for micrometers is very slim, for millimeters greater, for inches even greater yet, and for meters it is virtually guaranteed. The same goes for partitioning 'bald' by varying numbers of hairs or 'heap' by varying numbers of salt grains. We are much more likely to draw the same boundary for a Sorites in which we take away one-thousand hairs at a time than by taking away a single hair at a time.

What explains the sense of arbitrariness in (A)? A natural thought is that it must be the same thing accounting for the sense of indifference (or even irritation) in undergoing a Sorites paradox, namely that we do not use ordinary vague predicates in such a way to be able to draw such a boundary. It is difficult to imagine realistic situations in which drawing such a boundary would be of any practical value but rather simply a waste of time. Furthermore, the finer the partitioning of the domain, the less likely drawing such a boundary would be of any real use. It is useful to be able to distinguish human height, describe certain recognizable hair patterns, and distinguish larger piles from smaller ones. It is little wonder then that the degree of plausibility of Sorites paradoxes, ultimately based on fineness of partitioning, corresponds directly with the degree of informativeness and inversely with the usefulness of the vague expression involved. Measuring human height by micrometers is impractical but incredibly informative – nearly all individuals are distinguished. Contrarily, measuring height by whole meters is very practical (i.e. we can easily distinguish meters) but utterly uninformative.[34] Analogously, we do not measure baldness by single hairs – it is practically impossible. It would, however, be incredibly informative to be distinguished by number of hairs – too informative perhaps.

The kind of informativeness found in very plausible Sorites partitionings is just that, too informative. We would hardly ever have occasion to need this kind of information. And given the plausible assumption that natural languages evolve to suit the needs of their users, it seems rather unlikely that the semantics of our ordinary predicates would have evolved to make such fine-grained distinctions. In fact, this may very well explain why this ancient paradox has remained so obstinate: We are searching for objective boundaries that do not, and in fact *could* not, exist. The evolutionary considerations discussed earlier strongly suggest Wright's tolerance principle is a metalinguistic principle necessarily constraining natural languages. Moreover, we should recognize that a principle like WP makes possible the kind of semantic utility vagueness provides. For if slight differences actually did matter to the application of vague predicates, we would have little reason for thinking that speaking vaguely is most often sufficient for the purposes of ordinary conversation. What makes speaking vaguely efficient is our ability to supply enough relevant information without being too precise. And this, in turn, makes sense only under the assumption that either we have the capacity for speaking precisely and most often choose not to, or that sometimes we do not even have that capacity. The evolutionary arguments

[34] Warning: The usefulness of 'tall' increases with the unit chosen only to a point, and then becomes more and more impractical the greater it gets. We cannot casually observe miles, just as we cannot casually observe micrometers.

made earlier (or, put differently, Graff Fara's cost-benefit / efficiency considerations) establish the former, and obvious epistemic and conceptual constraints establish the latter.

Tappenden has argued that in uttering something like WP for corrective purposes, it is unnecessary that WP be true. Rather, it is sufficient that it be only not false.[35] If, as I suggest, instead the tolerance principle is a *metalinguistic* principle constraining all natural languages, we need not speak of its truth or falsity at all. We may think of tolerance as a metalinguistic imperative, and imperatives do not have truth-values. In essence, we have two seemingly competing accounts of what tolerance should look like:

(WP) For any tolerant predicate F and objects o_i and o_j, if o_i and o_j are observationally indistinguishable or saliently similar (along some relevant dimension), then Fo_i if, and only if, Fo_j.

(WPI) Classify casually indistinguishable or saliently similar objects (along some relevant dimension) alike.

WP simply formalizes the metalinguistic imperative WPI. We can think of tolerance this way because, rather than surveying all natural languages and verifying that in fact they all conform to WP, we have seen good reason for thinking that tolerance is a necessary constraint on all natural languages. Of course, we presuppose WP to hold, but this is simply because it formalizes the imperative WPI which itself constrains linguistic use. We can contrast this metalinguistic imperative with what intuitively verifies P1 and P2, namely descriptive facts of English. That Fo_1 and $\neg Fo_n$ are in fact true is *descriptive* in the sense that this is something a statistical survey of English use could reveal. Drawing this 'normative vs. descriptive' distinction does not solve the paradox, but does go some way toward explaining its infamous obstinacy.

This leads to a rather natural answer to the psychological question. We have seen good reason for thinking that tolerance is a necessary feature of natural language perceptual predicates.

[35] See Tappenden (1993).

After all, it seems crucial for understanding how vague predicates could in fact be useful. Further, to classify two objects indistinguishable in regard to *F* differently is to belie how the predicate is used. Also, if perceptual predicates are to be of any use, we must expect other speakers to be honest in their reports regarding those predicates. And this point generalizes to non-perceptual vague predicates. If ‘rich’, ‘adolescent’, and the like are to be of any use, we must assume that others are honest in their employment of those predicates. A significant part of this expectation is that speakers use these predicates in conformity with linguistic convention – recall Parikh’s observations concerning color predicates. In Grice’s terms, if communication is to be useful, it must be cooperative. And wasting others’ time with irrelevant, insignificant details is to opt out of cooperative communication. It is for this very reason that we immediately think that obvious truths like BIP, TIP, and HIP are undeniable. It is also for this very reason that undergoing an actual Sorites would be rather frustrating. In essence, we recognize that establishing such fine-grained boundaries frustrates the semantic utility of such predicates, namely expressing a sufficient amount of information without being overly precise. Put differently, and as Graff Fara maintains, it would be *inefficient* to draw such distinctions given ordinary (i.e. non-Soritical) conversational interests and purposes.

If this is right, the appropriate analysis of vagueness and the Sorites paradox falls squarely within the realm of pragmatics. As argued, the semantics of prototypical vague predicates depends intimately on features of conversational context. Furthermore, considerations of speakers’ expectations, interests, purposes, and the like are traditionally lumped into this category. My treatment operates within a broadly Gricean framework which sees communication as a subspecies of more general, rational cooperative behavior. I have suggested that an analysis of vagueness should track the assumptions, presuppositions, and other propositional attitudes constituting the common ground of particular conversations in which vague utterances take place. I argued that this is essential not only for accommodating what supervaluationists have termed *penumbral connections*, i.e. certain obvious inferences based on certain background assumptions, but also for understanding the constraints on the Sorites paradox (especially S3 and S4).[36] Employing this framework, I showed why we are wrong to suppose that inductive premises, independent of further background assumptions, are always plausible. Furthermore, by showing how abstracting from the (generally accepted) semantics of prototypical vague predicates to Sorites continuums is generally problematic, I have given good reason for thinking that the Sorites poses no interesting problems for the semantics of (at least) prototypical vague predicates. If this is right, I hope to have given sufficient reason for deflating the significance of the paradox. Finally, by placing the source of the paradox within a broader, evolutionary framework, I have offered a general solution consistent with but more encompassing than that of the contextualist (especially Graff Fara) and open to theorists of Vagueness from all camps. Vagueness is not usually a problem for ordinary language. We get by speaking vaguely, and if I am right, we could not help but speak vaguely. Rather, our fascination with this ancient paradox is entirely owing to the interesting problems it poses to set theory and classical logic. I hope to fill out this story in greater detail later with an eye to just this.

[36] See Fine (1975).

References

Stalnaker, Robert (1974) ‘Pragmatic Presuppositions’, *Semantics and Philosophy*, Munitz and Unger (ed.), New York: New York University Press.

Wright, Crispin (1975) ‘On the Coherence of Vague Predicates’, *Synthèse* 30, pp. 325-63.

Fine, Kit (1975) ‘Vagueness, Truth and Logic’, *Synthese* 30, pp. 265-300.

Unger, Peter (1979) ‘There Are No Ordinary Things’, *Synthese* 41, pp. 117-54.

Smith, Maynard (1982) *Evolution and the Theory of Games*, London: Cambridge University Press.

Grice, Paul (1989) *Studies in the Way of Words*, Cambridge, MA: Harvard University Press.

Kaplan, David (1989) ‘Demonstratives’, from *Themes from Kaplan*, New York: Oxford University Press.

Tappenden, Jamie (1993) ‘The Liar and Sorites Paradoxes: Toward a Unified Treatment’, *The Journal of Philosophy* 11, pp. 551-77.

Williamson, Timothy (1994) *Vagueness*, London: Routledge.

Raffman, Diana (1994) ‘Vagueness Without Paradox’, *Philosophical Review* 103, pp. 41-74.

Kamp, Hans and Barbara Partee (1995) ‘Prototype Theory and Compositionality’, *Cognition* 57, 129-91.

Parikh, Rohit (1996) ‘Vague Predicates and Language Games’, *Theoria* 27, pp. 97-107.

Clark, A. (1997) *Being There: Putting Brain, Body and World Together Again*, Cambridge, MA: Bradford Books / MIT Press.

Horgan, Terrence (1998) ‘The Trasvaluationist Conception of Vagueness’, *Monist* 81, pp. 316-33.

Soames, Scott (1999) *Understanding Truth*, New York: Oxford University Press.

Kennedy, Chris (1999) ‘Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison’, New York: Garland.

Graff Fara, Delia (2000) ‘Shifting Sands: An Interest-Relative Theory of Vagueness’, *Philosophical Topics* 28, pp. 45-81.

Stalnaker, Robert (2002) ‘Common Ground’, *Linguistics and Philosophy* 25, pp. 701-21.

Jucker, Smith, and Lüdge (2002) ‘Interactive Aspects of Vagueness in Conversation’, *Journal of Pragmatics* 35, pp. 1737-1769.

Heck, Richard (2003) ‘Semantic Accounts of Vagueness’, in *Liars and Heaps*, J.C. Beall (ed.), New York: Oxford University Press.

Thomasello, M. (2003) *Constructing a Language: A Usage-Based Theory of Language Acquisition*, Cambridge, MA: Harvard University Press.

Stanley, Jason (2003) ‘Context, Interest-Relativity, and the Sorites’, *Analysis* 63, pp. 269-80.

Shapiro, Stewart (2006) *Vagueness in Context*, New York: Oxford University Press.