

THE IMPLICATIONS OF VAGUENESS

by

by

Jesse L. Rester

Submitted to the

Office of Honors Programs and Academic Scholarships

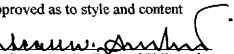
Texas A&M University

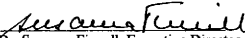
in partial fulfillment of the requirements for the

1998-99 University Undergraduate Fellows Program

June 15, 1999

Approved as to style and content


Dr. Scott Austin, Dept. of Philosophy


Dr. Susanna Finnell, Executive Director
Honors Programs and Academic Scholarships

Abstract

The Implications of Vagueness

Jesse L. Rester, Philosophy, Texas A&M University

1998-99 University Undergraduate Fellows Program, Dr. Scott Austin

If grains of sand are added, one by one, to a growing collection of sand on an otherwise empty table, there will eventually be a "heap" of sand on the table. It seems impossible, however, to specify the precise point at which the collection of grains becomes a heap. One grain of sand is certainly not a heap of sand. Does two grains comprise a heap? Can the collection be called a heap at three grains, at ten grains, at 500 grains? The commonly used term "heap" is *vague* – there is no clear line which demarcates the heaps from the non-heaps. The difficulty presented by this vagueness becomes clear when we examine the *sorites paradox*, a very old philosophical problem which is centered around the premise that the term "heap" has no precise definition. If it is impossible to specify exactly which objects are heaps and which objects are not heaps, how do we continue to use the term with such impunity? Is it possible for a system of logic to model the use of vague terms, if their application is often "neither true nor false," or, "only a matter of interpretation?" How are we to understand the role of vagueness within language?

Contents

- I. Introduction: The Sorites Paradox**
- II. The Theory of Supervaluations**
- III. Higher-Order Vagueness**
- IV. Many-Valued Logics and Degree Theory**
- V. Burgess' Proposal**
- VI. Scientific Realism and Nihilism**
- VII. Epistemology and Pragmatism**

Introduction: The Sorites Paradox

The paradox of the heap has been around for at least 2,000 years, probably originating in the work of Zeno of Elea and Eubulides of Miletus. In the ancient context the paradox figured prominently in debates between the skeptic and stoic schools of philosophy, and its use in public debate is recorded by Diogenes Laertius, Galen, and Cicero. As time moved on, however, the topic dropped out of view. Vagueness was apparently not considered an issue during the medieval era, and received only a cursory and intermittent recognition during the Renaissance and early Modern period. Recently, however, as issues of logic and language have again moved into the philosophical foreground, a sizable body of work has been devoted to the sorites paradox.¹ The perseverance of the paradox attests to both its philosophical import and to the lack of any simplistic solution.

Suppose that, in boredom one day, we begin to construct a heap-like collection of sand starting with a single grain and adding one grain at a time. Each time a grain is added we ask, "is the collection now a heap?" If this question is at first answered in the negative (two grains of sand do not comprise a heap), and eventually comes to be answered in the affirmative (this 100,000-grain collection is a heap), then it seems we must at some point specify a sharp division between heaps and non-heaps. The point at which we stop answering "no, the collection is not a heap" and begin to answer "yes, the collection is now a heap" indicates the location of this division. Unfortunately, the specification of this division between heaps and non-heaps seems completely arbitrary. If when the 76th grain of sand is added I proclaim, "it has now become a heap," you may very well disagree. It seems unlikely that either of us would be able to formulate a conclusive argument to support the truth of our position. Such an argument would need to appeal to the exact meaning of the term "heap," yet when we try to examine this meaning we encounter the problem of vagueness – "heap" does not seem to have a precise definition. If this is the case, there will be no specific point at which we can accurately, objectively *begin* to call our growing collection a

"heap," and thus it seems we can never say "yes, the collection is a heap," regardless of how many grains it contains.²

The most obviously contentious issue in this sorites reasoning is the claim that "heap" has no precise definition. "If a n -grain collection of sand cannot be arranged into a heap," we might say, "then a $(n+1)$ -grain collection of sand cannot be arranged into a heap either." This claim is called the "principle of tolerance" for vague applications; more generally, "if an instance is accurately described by the vague term in question then very similar instances are also accurately described by that term, and if an instance is not accurately described by the vague term in question then very similar instances are also not accurately described by that term." The principle of tolerance is the backbone of the sorites paradox and is often taken to be the defining characteristic of vague terms.

The principle of tolerance also allows "positive" versions of the sorites paradox to be constructed. We know that a 100,000-grain collection of sand can be arranged into a heap, so by the principle of tolerance a 99,999-grain collection can also be arranged into a heap, and again by the principle of tolerance a 99,998-grain collection can also be arranged into a heap, etc. The eventual result of this argument is that any collection of sand is a heap. In general, sorites reasoning takes as its premises the principle of tolerance and a single application, positive or negative, of a vague term. Step by step, grain by grain, the sorites reasoning leads us to universalize the application with which we began.

The issues associated with vagueness are far-reaching and apply to a great deal more than heaps. "Bald," "red," "adult," "tall," and "tadpole" are just a few prominent examples of vague terms. Suppose that we filmed the development of a young tadpole as it slowly changed into a frog. Beginning with the first frame of the film we ask, "is it a frog yet?" Frame by frame, we reach the conclusion that there are

¹ For a more detailed historical discussion see Williamson (1994). For historical source material see Keefe & Smith (1996).

² When a "collection of sand" is discussed in regards to the sorites paradox, it can be assumed that this collection is appropriately arranged to be as "heap-like" as possible. This allows us to concentrate our attention on a single axis along which "heap" is vaguely defined – in this classic example we are using the number of grains in the collection. "Heap" may be vague along other axes, such as the shape of the collection, but these other factors are unimportant for our discussion and should be taken as constant. More generally, any non-specified aspect of an instance used in sorites reasoning should be considered to be paradigmatic of the vague term in question.

no frogs. Any vague statement, which is to say the majority of human language,³ will be susceptible to similar lines of argument. For example, this sorites reasoning is prominently featured in "slippery slope" ethical dilemmas. Suppose that as a lobbyist it is acceptable that I give my old friend the senator a small, sentimental birthday gift worth only a few dollars. A one-cent price difference surely cannot make the difference in what is an "acceptable" gift.... The ethical problem of "where to draw the line" is seen to be another problem of vagueness. The sorites paradox, then, has significant implications for our understanding of language and of the things we do with language.

³ The extent to which vagueness pervades natural language will be discussed in Section VI.

Sorites Paradoxes

General Form:

- P1) The Principle of Tolerance: if (and only if) an instance is accurately described by the vague term, then very similar instances are also accurately described by the vague term.
- P2) Categorical Premise A: an instance which is accurately described by the vague term.
- P3) Categorical Premise B: an instance which is not accurately described by the vague term.

By recursively applying P1 to one of the categorical premises (using modus ponens), a conclusion is reached which directly contradicts the other categorical premise.

Example:

- P1) If an n -grained collection of sand can be arranged into a heap, then an $(n-1)$ -grain collection of sand can also be arranged into a heap.
- P2) A 100,000-grain collection of sand can be arranged into a heap.
- P3) A 2-grain collection of sand can not be arranged into a heap.
-
- 1) A 99,999-grain collection of sand can be arranged into a heap. (P1, P2, modus ponens)
- 2) A 99,998-grain collection of sand can be arranged into a heap. (P1, 1, modus ponens)
- 3) A 99,997-grain collection of sand can be arranged into a heap. (P1, 2, modus ponens)

Continuing this sequence, we arrive at the conclusion that a 2-grain collection of sand can be arranged into a heap, which contradicts P3.

Implications:

The possible responses to a Sorites argument are:

- 1) Reject P1
- 2) Accept the conclusion, and reject either P2 or P3.
- 3) Accept the contradiction.
- 4) Reject the validity of modus ponens.

The Theory of Supervaluations

Recent discussions of the implications of vagueness for logic can perhaps be said to begin with the interpretive distinctions seen in the theory of supervaluations.⁴ Non-vague predicates such as, "is an even number" can be understood in terms of positive and negative extensions. The positive extension of the predicate is the set of all particular instances to which the term can be accurately applied; in this case the set of all things which are in fact even numbers. The negative extension of the predicate is the set of all particular instances to which the term can not be accurately applied; in this case the set of all things which are not even numbers. For non-vague predicates, everything will fall into either the positive extension or the negative extension – it either is an even number or it is not an even number. In addition to these two extensions, however, the theory of supervaluations proposes that vague predicates have a third extension, or "penumbra." The penumbra is the set of instances which lie between the positive and negative extensions, the borderline cases where there is no simple fact of the matter as to whether the application of the vague predicate is correct.

Let us suppose that the vague predicate, "is a heap" has a penumbra which includes collections of sand between 10 and 10,000 grains. This means that a 9-grain collection of sand is definitely not a heap, and the statement, "a 9-grain collection of sand can be arranged into a heap" is definitely false. On the other hand, a 10,001-grain collection of sand definitely can be arranged into a heap. As for the penumbral instances which lie between these two paradigms, it is only a matter of interpretation as to whether they are heaps or not. According to the theory of supervaluations, the definition of "heap" does not include

⁴ Bertrand Russell makes use of these interpretive distinctions and anticipates many of their consequences in his seminal paper on vagueness in 1923. According to Keefe and Smith (1996, 1) the theory of supervaluations was first presented by Henryk Mehlberg in 1958, in a work entitled *The Reach of Science*. The term "supervaluations" was coined by Bas van Fraassen (1966), though he did not specifically apply the theory to the problems associated with vagueness. The details and applications of the theory of supervaluations were discussed and debated during the 1970's by a variety of people, including D.K. Lewis (1970), Michael Dummett (1975), J. Kamp (1975), Marian Przelecki (1976), D. H. Sanford (1976), and Kit Fine (1975). Fine's paper, "Vagueness, Truth, and Logic," is widely considered to be the most sophisticated and inclusive of these texts.

enough information to decide (for example) whether a 76-grain collection is a heap. Therefore the statement, "a 76-grain collection of sand can be arranged into a heap" is neither true nor false. The penumbra is variously considered to be a "truth-value gap," containing instances which fall between being truly described by the predicate and being falsely described by the predicate, or a "truth-value overlap," containing instances which can be said to be both truly *and* falsely described by the predicate. The penumbra is a deficiency in the meaning of the vague term, but the theory of supervaluations takes such deficiencies to be common and for the most part acceptable linguistic phenomena.

In order for supervaluations to provide a feasible description of language use, a few constraints need to be placed on the indeterminate truth value given to penumbral applications of vague predicates. For example, it does not make sense that a 76-grain collection can be called a heap while a similar 100-grain collection can be called a non-heap. Supervaluations uses the concept of the "sharpening" to avoid such incoherent uses of vague terms. The sharpening is a specific "precisification" of the vague term which eliminates the penumbra while preserving the truth value of any statements which were unproblematically true (or false) before the sharpening. The claim that, "a small collection of sand can be accurately called a heap while a similar, larger collection of sand can be accurately called a non-heap" is unproblematically false, therefore we are not allowed to sharpen the term "heap" in any way which could make this sentence true. One of many possible valid sharpenings for our heap example would be to consider collections of 76 or more grains to be heaps, and collections of 75 or less grains to be non-heaps. It is acceptable to use the term "heap" to mean "properly arranged collections of 76 or more grains," as long as it is understood that this is only one of the acceptable ways in which the term can be used. Proper use of vague language, then, should be described not in terms of individual applications but rather in terms of allowable sharpenings. At the higher level of interpretation, a sentence is said to be "supertrue" if and only if the sentence is true for all sharpenings, and a sentence is "superfalse" if and only if it is false for all sharpenings. Otherwise, the sentence cannot be categorized as true or false due to a deficiency (or perhaps overlap) of meaning. Since sharpenings are defined to be the possible precisifications which do not alter traditionally unproblematic truth evaluations, supertruth evaluations should not modify the basic,

accepted principles of classical logic. The theory of supervaluations claims that supertruth is equivalent to classical truth, and merely provides a more detailed explication of classical truth evaluations for use in the examination of problematic statements.

The principle of tolerance is one "problematic" statement specifically targeted by supervaluationism. It is perfectly acceptable, on the supervaluationist's account, to call a 76-grain collection a heap and a 75-grain collection a non-heap, even though these two objects are very similar. Furthermore, the principle of tolerance will be false for *all* possible sharpenings, and thus the sorites argument would seem to be based on a faulty premise. Any sharpening is a precisification, and will create a precise boundary somewhere, and wherever that line is drawn will be a counterexample to the principle of tolerance. According to the theory of supervaluations, the difficulty in agreeing "where to draw the line" is due not to the impossibility of declaring a precise boundary, but rather to the existence of many acceptable boundaries within the penumbra. By using the higher-level definition of truth as truth over all possible sharpenings, arguments over penumbral applications of vague terms can simply be shunted aside as "irresolvable due to the deficiency of the terms used." Unfortunately for supervaluationism, the theory now needs to explain how to distinguish the irresolvable arguments from the matters of fact; that is, we need some method for determining which applications of a particular vague term lie within the penumbra, and which applications are safely beyond this area of uncertainty. This "higher-level vagueness" leads eventually to the downfall of the supervaluationist project, but before we turn to this topic there is another aspect of the theory which should be observed.

Supervaluations leads to some rather unusual interpretations of logic and language. Some of these, including the rejection of the principle of tolerance, stand accused of violating intuition and practice. For example, Williamson (1994, 151-2) claims that the theory of supervaluations must reject at least four traditionally valid forms of argumentation (contraposition, the deduction theorem, or-elimination, and *reducto ad absurdum*) and that the theory thus "invalidates our natural mode of deductive thinking." Proponents of supervaluationism, on the other hand, have held that their interpretations are the more intuitive, and only revise areas of logic which were previously "problematic." We may agree with

Machina (1976, 178) that, "a battle of raw intuitions is liable to bore onlookers;" nonetheless, it should be beneficial to understand exactly what intuitions are at stake.

One of the most controversial claims of supervaluationism involves the treatment of disjunction.

The statement,

This object is a heap or this object is not a heap

is true for all sharpenings, and thus "supertrue." No matter where a sharpening draws its line, a particular object will always be on one side or the other. This is one example of supervaluationism affirming a traditionally true statement, which in more general form is called the law of excluded middle:

It is true that (P or not P).

Suppose that we are discussing a 76-grain collection of sand, an object within the proposed penumbra of "heap." That "this object is a heap or this object is not a heap," is, as we have said, true for all sharpenings. However, the claim that "this object is a heap" is false on some sharpenings, and likewise, the claim that "this object is not a heap" is false on some sharpenings. Neither of these later two statements are supertrue. Thus supervaluationism rejects what is called the law of bivalence:

(It is true that P) or (It is true that not P).

The disjunction (P or not P) can be evaluated as supertrue even if neither of the disjuncts (P), (not P) are themselves supertrue.

It is this critical interpretation which is the foundation of supervaluationism's rejection of the principle of tolerance. The following sentence is true for all sharpenings:

There is some number n within the penumbra of "heap," where a n-grain collection can be arranged into a heap, and a (n-1)-grain collection can not be arranged into a heap.

If n is replaced with *any* specific number, however, the result is a sentence which is *not* true for all sharpenings:

For a number N , a N -grain collection can be arranged into a heap, and
a $(N-1)$ -grained collection can not be arranged into a heap.

This sentence is supertrue for "some number n ," but it is not supertrue for any particular number.⁵

Remarkably, supervaluationism embraces and defends this position as intuitive. When speaking of a penumbral instance it would not be appropriate to assert that the object is undeniably, objectively a heap; nor would it be appropriate to assert that the object is undeniably, objectively a non-heap. Although we must of necessity make use of sharpenings in our day-to-day activities, from a more analytical perspective we generally admit that there is room for disagreement on such things. However, it *is* appropriate to firmly state that the object must be one way or the other, either a heap or not a heap. Fine (1975, 137-8) draws an analogy between this interpretation of vagueness and the semantics of ambiguity. Considered without context, the statement "X is a bank" is true only if X contains money *and* X contains a river (etc.) Assume that X contains a river but no money. Then the statement "X is a bank" is not definitely true, because it is uncertain what "bank" means here. Similarly, the statement "X is not a bank" is also not definitely true. However, the statement, "X is a bank or it is not a bank" *is* definitely true regardless of context. Both vagueness and ambiguity, according to Fine, lead us to rethink the traditional interpretation of the disjunction. Others remain critical, however, including Sanford (1976), Kamp(1981), and Sainsbury(1990).

The law of excluded middle and the law of bivalence are so similar as to be often confused, and indeed classical truth-functional logic cannot distinguish between the two.⁶ In proposing a significant

⁵ The relation of this issue to the previous discussion can be seen if the statement "X at some n " is considered to be equivalent to the statement, "X at n_1 , or X at n_2 , or X at n_3 ..."

⁶ If " Tp " stands for " p is true," then excluded middle is written as $T(p \text{ or } \neg p)$ while bivalence is $(Tp \text{ or } T\neg p)$. Traditional logic does not distinguish between " Tp " and " p " and thus the two statements are rendered identical. In supervaluational logic " Tp " is equivalent to " p for all sharpenings," and this allows excluded middle and bivalence to be differentiated.

differentiation between these two laws (accepting excluded middle and yet rejecting bivalence), supervaluationism attempts to understand and reconcile what Machina referred to as a battle of raw intuition over the logical treatment of the disjunction. More generally, this debate is symptomatic of a larger conflict of intuition concerning the reality and "depth" of vagueness within language. On the one hand, it does not seem appropriate or even possible to attribute precise boundaries to the definition and application of terms such as "heap." This intuition leads us to reject the traditional law of bivalence, as well as any other attempt to reduce vague language to precise extensions. On the other hand, we have seen that a lack of sharp boundaries around our concepts leads to instability and paradox, bringing us to the conclusion that our vague language is empty and inadequate to the task of demarcating any sort of concrete reality. Supervaluationism genuflects to these intuitions by allowing for "shades of gray" while yet confining the phenomenon of vagueness within a larger, overarching framework of precision (that is, a precisely defined positive extension, negative extension, and penumbra). Still, we will see that if supervaluationism is to resolve the sorites paradox, it must eventually confront these conflicted intuitions and find a justifiable, non-arbitrary way to locate the precise extensions of vague terms.


Higher-Order Vagueness

Supervaluationism has so far been presented in its most simplistic form, specifying three extensions for vague terms. Higher-order vagueness is the argument that the boundaries of these three categories must themselves be vague; that there are borderline borderline cases. In the original example, we proposed that nine grains can definitely not be formed into a heap, whereas ten grains is an arguable case. There appears to be little justification for choosing this particular cutoff point over other possibilities. Similarly, what justification could be given for specifying collections of 10,001 grains as definite heaps, while any lesser collection is arguable at best? In light of this argument, it appears as though supervaluationism has merely replaced one problematic boundary (between heaps and non-heaps) with two problematic boundaries. A particular division between definite heaps and arguable heaps seems, like the original sharp division between heaps and non-heaps, to be arbitrary and unreflective of the natural use of the term. However, if the boundaries of the penumbra are themselves irresolvably vague, then "in the positive extension" is itself a vague term, "all possible sharpenings" is a vague term, and supervaluationism will not be able to block the sorites paradox. (If a 10,001-grain collection is "in the positive extension of heap," then a 10,000-grain collection must also be "in the positive extension of heap"....)

The most obvious response to higher-order vagueness is to simply add higher layers of interpretation, replacing the sharp boundaries of the original three extensions with secondary penumbras. These secondary penumbras would contain the instances which are neither definitely within nor definitely *outside of the first-level penumbra*. Second-level penumbras will of course also create sharp boundaries, once again inviting the critique of higher-order vagueness, which can once again be countered by adding a third level of supervaluationist interpretation, etc. Returning to our heap example, a second level of supervaluationist interpretation would replace the 9-grain / 10-grain boundary with a secondary penumbra of, perhaps, 5 to 15 grains. Less than five grains would be "most definitely not a heap," and five to fifteen

grains would be "arguably a definite non-heap" (see diagram). A third level of interpretation would replace the 4-grain / 5-grain division with yet another penumbra, perhaps three to six grains. There are several possible endings to this procedure. Either the escalating levels of vagueness simply end somewhere, or perhaps they become insignificant to the understanding of the vague term, or else they should be thought of as continuing ad infinitum.

Regardless of the levels of interpretation invoked, the boundaries of the original positive and negative extensions remain a critical issue for supervaluationism, for it is these boundaries which distinguish the matters of fact from the (on some level) arguable cases. It is even questionable as to whether the other boundaries created by higher levels of supervaluationist interpretation are significant. For example, according to our proposed second-level interpretation of heap, a 6-grain collection is "arguably a definite non-heap." The statement that "a 6-grain collection can be made into a heap" is neither true nor false, and the 5 grain / 6 grain division is a valid sharpening. This seems nearly identical to saying that a 6-grain collection is within the first-level penumbra, and it is debatable as to whether all higher-level penumbras do not simply comprise a single, large penumbra. In any case, we need only focus on the critical boundaries of the positive and negative extensions. If these boundaries are vague, then "positive extension" remains a vague term, the sorites paradox remains unresolved, and our understanding of vague language remains indeterminate. If these boundaries are to be understood as naturally precise, their location will have to be justified in some non-arbitrary manner.



Many-Valued Logics and Degree Theory

The theory of supervaluations is only way in which non-traditional systems of logic have tried to adequately describe vague language. Other contenders include logical systems with more than two truth values, which have previously been developed for use in a variety of other applications. (For example, the statement "the roll of a fair six-sided die will land as a four" can conceivably be evaluated as one-sixth true.⁷) Like the theory of supervaluations, systems of many-valued logic have been applied to the understanding of vague predicates and as solutions to the sorites paradox. The most simple manner in which to do this is to invoke a three-valued logic where borderline statements take on an intermediate truth value known as "indefinite" or indeterminate." This framework is very similar to the three extensions of supervaluationism, but the resulting logic and semantics can vary a great deal. Another common approach is to use an infinitely-valued logic to model vague applications, assigning statements such as "this 76-grain collection is a heap" a value along the continuum between zero (completely false) and one (completely true). The fractional value assigned is intuitively interpreted as the "degree of truth" of the statement. In this approach an object's membership in the set of "heaps" becomes a matter of degree; the set of objects which are heaps has no distinct boundary and is called a "fuzzy set."⁸

Systems of many-valued logic vary a great deal in their treatment of semantics and logical argument. One important consideration is the determination of the truth of connective statements. We have seen that in the theory of supervaluations the truth of the "or" connective could at times act independently of the truth of the component statements, and similar possibilities are available to systems of many-valued logic. Most systems strive to generalize on the classical treatment of connectives (e.g. Machina 1976), but some

⁷ On the other hand, the statement "this roll of a die will land is a four" should only be evaluated as definitely true or definitely false. The critical difference is that the probable outcome of the roll of "a" die is a purely abstract, analytical matter, whereas predicting the roll of this particular die is a factual issue.

⁸ Rescher (1969) gives a survey of various many-valued logical systems and their applications. The concept of the "fuzzy set" was originally discussed by Zadeh (1965). I am indebted to Keefe and Smith(1996) for their extremely coherent discussion of the technical concerns involved in the development of a many-valued system of logic.

make a more radical departure (e.g. Edgington 1996). The most common treatment is based on the interpretation given by Lukasiewicz in 1930:⁹

value of (p and q) = minimum truth value among the component statements
value of (p or q) = maximum truth value among the component statements
value of (not p) = one minus the truth value of p
value of (if p then q) = one when the truth value of q is greater than p; otherwise,
one minus the truth value of p plus the truth value of q.

These definitions were intended for an infinite-valued system of logic, but they work quite well for three-valued systems if the intermediate truth value is considered as a degree of one-half.

Another item of contention among various theories of many-valued logic is the conditions that are required to make an argument valid. In classical logic validity is usually interpreted as the preservation of truth – an argument is valid if and only if it is impossible for its premises to all be true while its conclusion is false. Many-valued logics may maintain this interpretation, holding that an argument is valid only if when its premises have a truth value of one its conclusion also has a truth value of one (e.g. Tye 1994). Most theories of many-valued logic, however, prefer to invoke a more complex notion of validity. Perhaps an argument is valid only if its conclusion is always at least as true as the least true premise (Machina 1976), or perhaps only if the sum of the falseness of the premises is always greater than the falseness of the conclusion (Edgington 1996).

These diverse theories of many-valued logic provide a variety of slightly differing interpretations of vagueness and of the sorites paradox. Often the paradox is avoided by criticizing the principle of tolerance. The statement, "this 76-grain collection of sand is a heap" might be assigned a truth value of .316, and the statement, "this 75-grain collection of sand is a heap" assigned a truth value of .315. The principle of tolerance as applied to these two objects would be:

⁹ Keefe and Smith 1996, 37.

If this 76-grain collection is accurately described as a heap (truth value .316), then this 75-grain collection is also accurately described as a heap (truth value .315).

Due to the slight drop in truth value between the two component statements, this instantiation of the principle of tolerance will in most interpretations be something less than completely true. (The standard Lukasiewicz interpretation given above would assign it a truth value of .999.) Edgington (1996) claims that that the generalized principle of tolerance gains cumulative degrees of falseness from each of its partially false instantiations, resulting in a principle which is mostly or completely false. Even if the general principle is only partially false, its iterated use in the sorites argument could, on some interpretations of validity, lead to a false conclusion. Yet another way to avoid the paradox is to reject the validity of the modus ponens argumentation used in sorites reasoning. If we take as premises the instantiation of the principle of tolerance outlined above (truth value of .999) and the statement "this 76-grain collection is accurately described as a heap" (truth value of .316), then the modus ponens conclusion, "this 75-grain collection is accurately described as a heap," has a lower truth value (.315) than either of the premises. On some interpretations (e.g. Machina) this means that modus ponens is not a universally valid form of argument.

The internal debates among supervaluationism and the various forms of many-valued logics are both fascinating and complex. There seems to be no end to the discussion over how best to model intuition and practice: three logical values or an infinity of logical values, how should we treat the "or" connective, is a "degree of truth" even comprehensible¹⁰, how should validity be described, etc. On a broader scale, however, all of these logical theories encounter similar problems in addressing higher-order vagueness and the "depth" of vagueness within language. The extensions created by a three-valued logic encounter the same difficulty as the three extensions of supervaluationism; a problematic division between "true

¹⁰ See Haack 1980, and for a defense, Sainsbury 1990, 43-4.

heap" and "false heap" is simply replaced with two problematic divisions in the true heap / indefinite / false heap schema. For infinite-valued logics, this same problem takes the form of deciding what degree of truth to assign to a particular statement. For example, should the statement, "this 76-grain collection is a heap" be assigned a truth value of .315 or a truth value of .316, or should it really be .378? Where should we locate the critical division between truth value 1 (completely true), and all lesser degrees of truth? A strong intuition against the arbitrariness and inadequacy of such precise description must be confronted. As we saw with supervaluationism, it is the actual application of logical theory to vague natural language which is most difficult – the attempt to assign precise values of truth to the use of terms such as "heap." This difficulty takes different forms in different theories, whether it is the attempt to locate the problematic boundary of the penumbra, the attempt to locate the similar boundary between a truth value of one and some lesser degree of truth, or the attempt to decide exactly what degree of truth to assign to a particular statement. The project of describing vagueness within an overarchingly precise framework seems to unfaillingly run aground against our intuition toward tolerance.

Burgess' Proposal

One attempt to locate and justify precise extensions for vague predicates is made by J. A. Burgess (1990). An examination of Burgess' proposal will help to clarify the difficulties involved with applying precise logical frameworks to vague natural language. Burgess uses a basic supervaluationist framework (positive extension, negative extension, and penumbra), defining the borders of the extensions as an empirical characteristic of language-use among a population. The method proposed to determine the precise extensions of a vague statement is to ask a group of "competent speakers" to identify the instances to which the vague statement definitely applies. The set of instances to which most of these people would unhesitatingly apply the vague statement is the positive extension of the statement, for that group, at that time. A similar procedure could be carried out to identify the negative extension, and the remaining instances would fall into the penumbra.

Unfortunately, this simple formulation is itself full of vague terms, such as "most" and "competent speaker." Before we could apply Burgess' method to any vague term, including the term "most," we have to know what the definition of "most" is. Thus the terms used to formulate the method could only be made precise by appealing to some higher, as yet unspecified, justification. If the terms used to formulate the method remain irresolvably vague, the sorites paradox can be used to undermine the coherence of those terms, thereby undermining the formulation of the theory and its ability to interpret language. This difficulty is yet again an aspect of higher-order vagueness, which we first encountered as vagueness within the terms "positive extension" and "negative extension." Higher-order vagueness, then, can be more broadly understood as "vagueness within the metalanguage," where the "metalanguage" is the language of the theory with which other language is described.

Burgess' strategy in responding to higher-order vagueness is to give a separate, higher-level analysis of each vague term within his metalanguage. For example, agreement among one-half or less of the population is "definitely not most," agreement among all but one is "definitely most," and all other

possibilities are "arguably most." Now that "most" has been given three precise extensions it can be used as part of a precise metalanguage. It is interesting that such precisification of Burgess' vague metalanguage turns what seems to be a reasonable, common-sense interpretation into an arbitrary and inadequate reduction of natural language. According to the "precisified" method, we can apply the term "heap" with confidence if all but one of us agree – yet if there are two dissenters the application of the term becomes merely a matter of interpretation.¹¹ Burgess' precisification of his metalanguage is unjustified and arbitrary, and therefore his interpretation of all other vague terms becomes unjustified and arbitrary, an imposition upon (rather than a reflection of) natural linguistic practice.

A central issue for theories of vagueness, then, is whether the metalanguage is to be vague, reflecting an intuition toward tolerance and the natural depth of vagueness within our language, or whether the metalanguage is to be made precise, reflecting the need to avoid indeterminacy and paradox. Without a non-arbitrary method of mapping precise extensions to natural language, attempts to formulate precise metalanguages seem inevitably inadequate to the task of describing those natural languages. Since such a method has yet to be found, we are forced to consider the precisification of natural languages to be a process of redefinition, not of description. If vagueness cannot be "redeemed,"¹² we must either reject vague language as incoherent, or else we must find some way to deal with linguistic indeterminacy, ignorance, or paradox from within a fundamentally vague interpretation. The following sections will expand upon these possibilities.

¹¹ Burgess should have at least taken into account population size, as agreement of all but one in a population of 10 million is vastly different than agreement of all but one in a population of 10. Even with this refinement, however, it is strongly counter-intuitive that the term "most" could be non-arbitrarily reduced to a specific percentage of the population.

¹² That is, if vagueness cannot be described precisely at a higher level of interpretation, if vagueness cannot be adequately modeled with a precise metalanguage.

Scientific Realism and Nihilism

The sorites paradox and related problems of vagueness have driven many theorists to the conclusion that any determinate and non-paradoxical knowledge must be precisely formulated. Bertrand Russell, for example, writes that most of the problems of formulating knowledge could be most adequately addressed from within the physical sciences, and that "science is perpetually trying to substitute more precise beliefs for vague ones."¹³ Once we accept that vagueness is a deep, pervasive, and yet paradoxical aspect of our language, one of the first inclinations of the scientifically-oriented theorist is to simply remove such flawed concepts from consideration.

By way of example, and to elucidate this eliminativist position toward vagueness, we can take a detour into philosophy of mind where the traditional "mentalist" perspective has recently been confronted with an "advancing tide of neuroscience." In this arena theorists such as Paul and Patricia Churchland have strongly pressed the idea that traditional "folk psychology" – the vague concepts of desire, fear, belief, self, intention, etc. – will soon be dismissed and replaced by some version of "the vocabulary required to characterize our kinematical states."¹⁴ Paul Churchland gives a rather disconcerting example of this new vocabulary when he asks us to consider the grafting of transducers into the brain, "to convert a symphony of neural activity into (say) microwaves radiated from an aerial in the forehead" for the purpose of communication.¹⁵ Rather than working with vague linguistic concepts, these brains would be able to directly relate precise "kinematical states." This potential improvement on the human condition could be considered a demonstration of the inadequacy of vague language as a characterization of actual physical interaction between individuals. For the eliminativist vague language is simply inarticulate,

¹³ Russell (1923, 12). Although the historical significance of Russell's paper is undeniable, his analysis suffers from a confusion between vagueness and ambiguity. Vague terms have borderline cases, whereas ambiguous terms have more than one meaning (any one of which could be either precise or vague). The notion of using a "less vague" term is not really applicable; a term either is, or is not, precisely defined (refer to the discussion of vague boundaries at the end of section III). On the other hand, a "less ambiguous" term is simply one which has fewer meanings. For a discussion of this aspect of Russell's paper see Williamson (1994, 64-69).

¹⁴ Churchland (1981, 222).

comprising a "flawed theory" which is incapable of describing concrete, material reality. Furthermore, most eliminativists argue that vague terms such as "self" cannot be reduced to more precise understandings, because (surprise) no underlying precise correlates can be found. A brain wired for direct communication would understand other individual brains, "in roughly the same fashion that your right hemisphere 'understands' and 'conceives' of your left hemisphere – intimately and efficiently, but *not* propositionally."¹⁶

In the physical model, of course, the two halves of the brain do not "understand" or "conceive" of each other at all, they simply participate in a physical interaction. Eliminativism leaves no room for the vagaries of "internal events" or "abstract conceptualization." A common short criticism of Churchland is that his theory has forgotten that it *is* a theory, a heuristic "understanding" rather than a necessary truth. Since theories are propositional (that is, they are abstract conceptualizations to which a self is related through belief, doubt, or other "psychological attitudes"), the context of Churchland's theory seems to contradict it's content. The physical/behaviorist model does appear to create something of a methodological problem, a certain *imposition* of theory, but it is not so easy to nail down. Churchland's response to this criticism is that considering theories to be propositions begs the question in the competition between his own theory and the theory of folk psychology. If Churchland is right, we need to take the meaning of linguistic constructs, along with everything else, to be an aspect of material interaction. We would do better to formulate criticism along empirical lines, to say that the precise physical model is simply insufficient to describe experience. Unfortunately for Churchland, the existence of "internality" and the associated propositional attitudes is not a pure theoretical *assumption*, it is a qualitative reflection of experience. Neuroscience may enjoy some measure of predictive success, but as a precise model it is not adequate to characterize subjective perception, reflectivity, and conceptualization. Certainly scientific reflection can *expand* upon our experience, but if it is to be empirical it should not undermine or reject the very existence of that experience. The eliminativist

¹⁵ Churchland (1981, 221).

position cannot account for the internality of experience, so it is scientific realism, rather than "folk psychology," which is the inadequate theory. Churchland should be expected to press the point, and to criticize the portrayal of (for example) "self" as an immediate aspect of experience. Surely it is difficult to claim that the concepts of folk psychology are contained *within* non-reflective experience. Rather, by claiming that neuroscience is inadequate to describe the "internal events" of experience, what is meant is that "internal event" is an intuitively *better* characterization than is neuroscience. Some significance might be attributed to the presumed goal of an adequate "characterization" of experience rather than the "prediction" of behavior, but this presents a false dichotomy. What is being predicted if not future experience?

The upshot of this discussion is that vague language yields a better description of experience than do precise frameworks of understanding. "I stepped on a heap" is a more adequate, more reflective characterization of this occurrence than is "molecule A exerted a force on molecule B" (or whatever). It would no doubt give eliminativists fits to see their view criticized by such a vague, undefined term as "better," referring to an irreducible, qualitative judgement on the ability of a linguistic structure to reflect experience. Churchland writes that, "an introspective judgement is just an instance of an acquired habit of conceptual response."¹⁷ Perhaps it is the lack of an intimate familiarity with a "completed" framework of neuroscience which leaves us with lazy "conceptual habits." On the other hand, perhaps it is the impracticality and awkwardness of the precise model which prevents it from being used in all but the most esoteric of situations. The methodological issue which arises is whether theory should be primarily accountable to the vagaries of experience (as "introspective judgement"), or whether it should instead be accountable to other theory (as "an acquired habit of conceptual response"). More properly, is a theoretical reflection of experience most adequately characterized in practical terms such as "judgement," or in the quantitative dictates of lab results?

¹⁶ Churchland (1981, 221).

¹⁷ Churchland (1981, 208).

It is interesting to note that precisely defined statements can be evaluated without appeal to experience. For example, if "heap of sand" is arbitrarily defined as any properly arranged collection of 76 or more grains, and "this thing" is defined as a properly arranged collection of 78 grains, then nothing could be more straightforward than the evaluation of the claim, "this thing is a heap." It is the broad application of such precisified frameworks to experience which becomes problematic. Suppose that we arrived at a completely precise definition of heap, including specified tolerances, a geometric description of "properly arranged," and a molecular description of "grain" based on size, composition, cohesiveness etc. The attempt to find (or even to create) a referent for this precise concept in experience will encounter the practical difficulties caused by reference to unobservables, underdetermination,¹⁸ and all the problems which continue to plague the attempt to characterize experience in purely precise terms. To take a simplified case, consider a precise "heap" of 76 or more grains. In order to ascertain whether an experienced "thing" is a heap, we have to count the grains. What if we miscount? Perhaps we decide it has 77 grains, but a second count reveals that it has only 74. Were some grains misplaced? How could we know whether what we experienced was a (precisely defined) heap or not? It is impossible to precisely characterize experience with any degree of certainty.

The foremost inadequacy of precise models, then, is that taken alone they have little "practical character;" they do not "work" unless they are integrated with more colloquial (i.e. imprecise) understandings.¹⁹ It is this unavoidable vagueness of experience and practice which thwarts the eliminativist project. Since immediate experience must be the starting point of any investigation, and

¹⁸ Reference to unobservables means that the characterization of experience in precise terms ultimately appeals to concepts such as "molecular bond" which (to the extent that they are precise) are only inferred from experience, not directly observed. This leads to the related problem of underdeterminism, which is that a variety of precise frameworks could potentially be inferred from the same observations (for example, wave and particle models of light). The *realist* claim that an instance of a precise "heap" has been found is thus clearly seen to be an misplaced imposition of a particular theoretical framework, similar to the claim that, "I saw a wavelength."

¹⁹ For example, in the employment of the "applied sciences" such as medicine or engineering, the precise frameworks of chemistry and physics take on a certain practical character within the larger stream of experience. On a broader level, practice / theory divergence can be seen within the theoretical sciences themselves, as idealized, precise models come into conflict with the actual contextual exigencies of scientific practice. This line of critical argumentation is due, of course, to the work of John Dewey and other Pragmatist philosophers.

because it is not possible to adequately and completely characterize this experience with precise conceptual frameworks, we must temper the ideal of precision and acknowledge the implications of vagueness. So far, however, we have found no adequate assessment of the sorites paradox. If the inclusion of vagueness within our conceptual framework allows for a better characterization of experience, and if this vague language is incoherent, then we are in distinct epistemological trouble. For some theorists (e.g. Unger 1979), a consideration of the sorites paradox leads ultimately to nihilism. Also called "despair" and "intellectual suicide,"²⁰ the nihilist position should be viewed as a product of the failure of eliminativism. Vagueness cannot be avoided and cannot be reduced to an underlying precise level of description. However, an acknowledgement of the depth and irreducibility of vagueness is only one aspect of our conflicted intuition. Nihilism does not take into account that vagueness *does* have "practical character," that vague language holds together, provides for at least a limited sense of stability and communication, and allows us to get through the day without bumping into tables and tripping over heaps. Therefore nihilistic despair can be seen as still holding to the false ideal that meaning must be precise (that vagueness must be eliminated) if language is to be coherent. How can we reconcile the functionality of vague language with paradox? Could there be some way of addressing paradox from within a fundamentally vague interpretation?

²⁰ Williamson (1994, 165).

Epistemology and Pragmatism

One move which has been fairly successful at resolving the sorites paradox while still admitting to the depth of vagueness is the interpretation of vagueness as ignorance. If we cannot discover a non-arbitrary way of locating a precise border between heaps and non-heaps, it does not immediately follow that no such border exists. Perhaps language does create such a boundary, but for vague terms we simply do not know where the line is located. Perhaps our inability to characterize experience precisely does not have significant ontological implications, but is a mere reflection of inadequacies in our (vaguely defined) perceptual and cognitive mechanisms.²¹ In this epistemological interpretation of vagueness, the term "heap" does in fact have a precise meaning. It might even be, "a properly shaped collection of 76 or more grains." The term is vague, however, in that we are incapable of discovering that precise meaning. Experience is most adequately represented with vague language because we are incapable of observing precise ontological boundaries – as was discussed earlier we cannot with certainty discriminate a 76 grain collection from a 77 grain collection, even if we counted.

One of the foremost criticisms of the epistemic interpretation of vagueness is that it yields an unacceptable externalization of the meaning of our language. Of course, it is reasonable to hold that particular uses of a language are not purely intentional, but are also influenced by external factors such as traditional meaning. However, the epistemological view seems to imply that the users of a language, even taken as a whole, have only a partial understanding of what we say. If it is impossible to specify the boundaries of meaning for a term such as "heap," then claiming that such boundaries exist seems to force us to claim that it is impossible to know what "heap" means. In defense, epistemic theorists reject the

²¹ To follow an example from the previous section, perhaps if our eyes were more sensitive we would be able to actually see and experience a particular photon (or some precise aspect) of light. If this were the case, we could indeed specify precise experience. Of course our perceptions do *not* provide that kind of experience, so we are "trapped" within vague reference and must ultimately disclaim any discussion of precise physical exigencies. We need further justification to suppose that reality might be precise *behind* our experience. The epistemic theory attempts to provide such a justification.

identification of "understanding" with verification, and argue that it is not necessary to "specify" something in order to comprehend meaning. As Timothy Williamson explains, "the measure of full understanding is not possession of a complete set of metaphysically necessary truths but complete induction into a practice... a practice that does in fact determine a meaning."²² Even if we may not be able to explicitly specify the meaning of a linguistic practice, we can still say that the meaning is determined by the practice. Linguistic participation is a form of understanding, albeit a vague understanding; in the epistemic view participation is the most direct and the best possible understanding of vagueness.

Although we are "trapped" within vague ignorance, epistemic theorists claim we can deduce that the meaning of vague terms must contain unknown precise boundaries. The sorites paradox is taken to be an argument against the existence of tolerance – vague terms must really be bivalent, or else they would be incoherent.²³ Such deduction allows us to know that precise boundaries must exist, even if it is impossible for us to discover where they are. Therefore the epistemic view finally allows us to salvage conflicted intuition, since it accounts for both the depth and the functionality of vague language. Though tenable, the epistemic theory remains a very *negative* characterization of vagueness. Paradox and indeterminism have been resolved, but we are left with ignorance,²⁴ and we must conclude that precise frameworks have little place in practice and experience. Vagueness is truly "deep," but it is existential rather than ontological, and our understanding continues to gaze (with unrequited longing) toward an externalizing, precise ideal.

A similar, but more affirmative response would allow vagueness to be both *paradoxical and* functional, would yield an instrumentalist assimilation of precise frameworks, and would provide a less alienating interpretation of experience. To William James, for example, words taken "at face value" are

²² Williamson (1994, 211).

²³ Referring to page 6, the argument would be taken as a *reductio ad absurdum* of P1. Williamson (1994, 187-98) gives further arguments, in addition to the sorites, to explain why vague terms must be "empty" and incoherent if vagueness is taken to exclude bivalence.

²⁴ It is difficult to judge the expansiveness of this ignorance, but it may turn out to be global where vague terms are concerned. Certainly precise characterization could not be assigned to "borderline cases," but higher-order issues must also be addressed. Presumably we are also ignorant of the exact boundaries of our knowledge...

simply inadequate, and yet language remains as a "provocation" and a "richness."²⁵ A pragmatic affirmation of vagueness takes paradox and indeterminacy as indicative of a fundamental openness within the ontological process, rather than as a gulf between experience and the real. In this view theories of logic and definition can be embraced as useful in certain situations, as long as we do not take the assignment of precise values too seriously (we have found that it is impossible to take them seriously anyway). And finally, experience and participation remain as the ultimate context of reflection. Perhaps vague language is more adequate *because* it is "empty" and paradoxical.

In summary, vagueness is a deep and significant aspect of our language. A specific reduction of vague language to an underlying precision is unjustifiable, so it is impossible to model vagueness with a precise logic or metalanguage. Some theorists believe that we should eliminate vague language as inarticulate, but this is also impossible because precise frameworks (by themselves) are insufficient reflections of experience. The failure of the ideal of precision, and the paradoxical aspects of vagueness, could potentially drive us toward nihilism or to a characterization of vagueness as ignorance. However, an acknowledgement of the transparency of vague language allows for an instrumentalist interpretation of precise concepts and an affirmation of experience.

²⁵ Gavin (1992, 69-76).

References

The following is a list of works discussed in this paper. For a more complete bibliography on vagueness, see Keeffe and Smith (1996, 338-45) or Williamson (1994, 307-19).

Burgess, J. A. 1990. "The sorites paradox and higher-order vagueness." *Synthese* 85: 417-74.

Churchland, P.M. 1981. "Eliminative Materialism and the Propositional Attitudes." In *Mind and Cognition*. 1996. Ed. William G. Lycan. Blackwell, Cambridge. 206-23.

Dummett, M. 1975. "Wang's paradox." In Keeffe and Smith, 99-118.

Edgington, D. 1996. "Vagueness by degrees." In Keeffe and Smith, 294-316.

Fine, K. 1975. "Vagueness, truth, and logic." In Keeffe and Smith, 119-50.

Gavin, W. 1992. *William James and the Reinstatement of the Vague*. Temple University Press, Philadelphia.

Haak, S. 1980. "Is truth flat or bumpy?" In *Prospects for Pragmatism*, ed. D. Mellor, 1-20. Cambridge University Press, Cambridge.

Kamp, J. 1975. "Two theories about adjectives." In *Formal Semantics of Natural Language*, ed. E. L. Keenan, 123-55. Cambridge University Press, Cambridge.

Kamp, J. 1981. "The paradox of the heap." In *Aspects of Philosophical Logic*, ed. U. Monnich, 225-77. Reidel Publishing, Boston.

Keeffe and Smith, eds. 1996. *Vagueness: a Reader*. MIT Press, Cambridge, MA.

Lewis, D. K. 1970. "General Semantics." *Synthese* 22: 18-67.

Machina, K. F. 1976. "Truth, belief, and vagueness." In Keeffe and Smith, 174-203.

- Mehlberg, H. 1958. *The Reach of Science* (excerpt). In Keefe and Smith, 85-88.
- Przelecki, M. 1976. "Fuzziness as multiplicity." *Erkenntnis* 10: 371-80.
- Rescher, N. 1969. *Many-valued logic*. McGraw-Hill, New York.
- Russell, B. 1923. "Vagueness." In Keefe and Smith, 61-8.
- Sainsbury, R.M. 1987. *Paradoxes*. Cambridge University Press, Cambridge.
- Sainsbury, R. M. 1990. "Concepts without boundaries." In Keefe and Smith, 251-64.
- Sanford, D.H. 1976. "Competing semantics of vagueness." In Keefe and Smith, 281-93.
- Tye, M. 1994. "Sorites paradoxes and the semantics of vagueness." In Keefe and Smith, 281-93.
- Unger, P. 1979. "There are no ordinary things." *Synthese* 41: 117-54.
- van Fraassen, B. 1966. "Singular terms, truth-value gaps, and free logic." *Journal of Philosophy* 63: 481-95.
- Williamson, T. 1994. *Vagueness*. Routledge, London.
- Zadeh, L. A. 1965. "Fuzzy sets." *Information and Control* 8: 338-53.