VAGUE PREDICATES AND LANGUAGE GAMES

Rohit PARIKH*

* City University of New York, Dept. of Computer Science, Box 750, Graduate Center, 33W 42nd Street, New York, NY 10036-8099, USA[1].

BIBLID [ISSN 0495-4548 (1996) Vol. 11: No 27; p. 97-107]

ABSTRACT: Attempts to give a Logic or Semantics for vague predicates and to defuse the Sorites paradoxes have been largely a failure. We point out yet *another* problem with these predicates which has not been remarked on before, namely that different people do and must use these predicates in individually different ways. Thus even if there *were* a semantics for vague predicates, people would not be able to *share* it. To explain the occurrence nonetheless of these troublesome predicates in language, we propose a different approach based on asking the question, "How do these vague predicates help people to communicate with each other?" We show that in general, even though different people assign different extensions to vague predicates, they usually benefit from receiving information framed in terms of them.

Keywords: Vagueness, Sorites paradoxes, Fuzzy Logic, Super-valuations, Mutual Knowledge, Utility, Nash Equilibrium.

Suppose I look at my watch and announce to my friend "It is now 4 PM". His response will probably not be, "I am glad to know that" or "Thanks for telling me," but more likely, "Why are you saying that? Is something happening at 4 PM?" While we all want to learn facts, we, or at least those of us who are not trivia buffs, do not learn them just to pass the time. There has to be some application, and probably some decision, immediate or eventual, that depends on hearing this sentence.

In other words, language is embedded in practice –we *act*, on our beliefs and we are diligent in correcting those beliefs that are important in our decision making. Still, we like to think that the *truth* of our beliefs is a separate concern from their part in our activities of running our lives and making decisions. We cannot avoid having beliefs that are not immediately useful, but we do hope that they are all true, and we take it for granted that being true is a property of beliefs that is independent of how or if they influence our practices. Thus, I believe that Toledo is in Ohio, and though I do not plan to make any use of this belief, I do assume that the truth of it is quite independent of that fact.

One purpose of this paper is to argue that this separation between truth and utility is not sustainable in the domain of the so-called vague predicates, even though it would be nice to have it, and though we *can* perhaps have it elsewhere.

I shall argue instead for an anti-representational view of vague predicates. Such a view of language has been espoused by Dewey, the later Wittgenstein, and Rorty, and emphasizes assertibility rather than truth. To quote Robert Brandom,[2]

> There is an approach to language shared by Dewey and the later Wittgenstein, which attributes little or no importance to the notion of truth. According to this view, language is best thought of as a set of social practices.

And again,

> One of the striking features of the later Wittgenstein's thought is the replacement of his earlier question 'What are the facts?' by the question 'What are we entitled to say?'[3]

Why does logic fail us?

One reason that a concept of truth does not work with vague predicates is the sorites paradoxes. To give one example, it is possible to construct a series of colored patches p_i such that (i) the first patch p_1 is (definitely) red, (ii) the last patch p_n is not red, and (iii) patch p_i is indistinguishable by human eye from patch p_{i+1}. A consequence of (iii) for most of us would be (iiia) patch p_i is red iff patch p_{i+1} is red.[4] (i), (ii) and (iiia) lead at once to a contradiction.

There have been various suggestions in the literature oriented to getting us out of this predicament. Briefly, the proposals are:

A) *Vague predicates are incoherent.* Hence no man is bald or tall and there are no heaps. Thus the paradox is avoided by denying i). This is an approach taken by Wheeler[5] and Unger[6] and while it solves our logical problems, it is surely too radical for most of us.

B) *The epistemic approach.* This approach denies (iiia) and allows that there is an i such that patch p_i is red and p_{i+1} is not and we just do not know which i this is. This approach is counter-intuitive, but it also solves the problem.

C) *Fuzzy Logic.* Under this approach, due largely to Zadeh,[7] there are degrees of redness and as i increases, this degree decreases so that with $i = 1$ this degree is 1 and with $i = n$ the degree is 0. However, fuzzy logic has difficulty with truth functions. In particular, if patch p_i has degree of redness equal to .5, then the statements $Red(p_i) \wedge \neg\ Red(p_i)$ and $Red(p_i) \vee \neg\ Red(p_i)$ both have degree of truth equal to .5, surely a counter-intuitive result.

D) *Super-valuations:* Under this approach,[8] there are many specifications for a vague predicate. Here a specification is a partial assignment of truth values to statements of the kind "this is red", "this is bald", etc, and one specification S extends another specification S', if every statement that is assigned a truth value by S' is also assigned a truth value by S, and moreover the same one. Kit Fine (1975) assumes that every specification has a complete, admissible extension, though not every complete extension may be admissible.

The super-truth theory then goes: "A vague sentence is (super) true if it is true for all admissible and complete specifications". However, it is implausible that there is such a collection of admissible specifications and moreover, the theory yields the consequence that there will be two indistinguishable patches p_i, p_{i+1} such that the statement $Red\ p_i$ will be super-true and $Red(p_{i+1})$ will not be.

E) *Non-classical Logics*: One possibility, suggested by Hilary Putnam,[9] is to deny $(\forall i)\ (Red(p_i) \rightarrow Red(p_{i+1}))$ but to use intuitionistic logic so that the consequence $(\exists i)\ (Red(p_i) \wedge \neg\ Red(p_{i+1}))$ does not follow. Other suggestions include using Kleene's 3-valued logic. However, any logic which accepts (iiia) will run into difficulties since all we need, to go from (i) to the negation of (ii), is modus ponens and even the weakest logic will allow this transition.[10]

Finally, even if we are prepared to give up logic, we encounter the same difficulty with a simple-minded notion of assertibility. Since "Patch p_i is red" is assertible and "Patch p_n is red" is not assertible, there will be a last $i < n$ such that "Patch p_i is red" is assertible. It follows that "Patch p_{i+1} is red" is not assertible. But how can we possibly explain a set of social practices which demands, of two indistinguishable patches, that we *may* say of one that it is red and of the other that we *may not* say that it is red?

In as much as assertibility is treated as a predicate, it is going to be subject to the same problems as truth, indeed more so. For we *can* contemplate that there is an i, rather like Gödel's undecidable sentence, such that "Patch p_i is red" is true and "Patch p_{i+1} is red" is not true, though *we* cannot tell which i it is. But this cannot hold for the notion of assertibility according to any plausible social practices.

So our situation is that we have a number of different approaches, all of which work, more or less and none of which is really satisfactory. I would like to propose that we start by asking three questions: "What are the contexts in which we do use the vague predicates?", "What properties of vague predicates are needed for them to be useful in these contexts?", and "What limitations of vague predicates prevent us from using them in other contexts in which we do *not*, normally use them?"

In other words, I propose that we start by exploring the language games that we can usefully play with vague predicates, and given that they *can*, play a role in certain language games, ask in what way is this role more restricted than the role that can be played by a truly bivalent predicate.

The rest of this paper will devote itself to these questions.

Learning Social Practices

Clearly, if participating as a competent speaker in a language requires familiarity with linguistic practices, then such practices (to the extent that they are not inborn) must be learned. And as I have argued elsewhere[11] not only can we not assign truth values to vague predicates, they also cannot be learned in a way that is consistent across a community. In other words, it is unavoidable that different speakers will use the same vague predicate differently, and that even the same speaker will use the predicate differently on different occasions.

There are two theoretical reasons why this must happen. One is that a predicate like "red" or "bald" can be applied to innumerable things. Since many of these predicates lack definitions, they tend to be learned by ostension. But then it is impossible that we could learn the proper usage of these words (if there *were* such a proper usage) from other people for *every* possible case. Rather, we learn the usage of

these words in some few cases and then we extrapolate. These extrapolations will agree to some extent since we humans are rather similar to each other. But we are not exactly similar. There are important differences between any two human beings, which will lead to differences in the way in which they extrapolate the same data.

The perception of colors, for instance, depends on certain non-repeatable physiological processes. If I look at a red spot for a few moments and then at a white wall, I will see on the wall a green spot that is not actually there. The appearance of the spot is caused by the fact that the red receptors in a certain part of my retina are tired, and the color green, which is complementary to red, is emphasized over the corresponding visual region. Since what we see is always somewhat affected by what we saw just before, our judgments will also be affected. Thus, for purely theoretical reasons we should expect *some* difference, small or large, in the assignment of color words to objects by different people. And our question now is whether this does actually happen.

It should be enough to cite one example. A set of colored squares was projected onto a screen[12] and the audience was asked to count the number of red squares and the number of blue squares. In one experiment, the number of red squares reported ranged from 12 to 30; the number of blue squares reported ranged from 0 to 27. In other words, members of the same audience issued different reports about the same overhead transparency.

Admittedy, the squares were chosen so as to bring out any latent differences. But also, these differences *were* there to be brought out and were substantial. It is also clear that among the projected squares there were some for which many of the audience would say that they were red and many others in the same audience that they were not. This fuzziness of the linguistic practice poses a challenge to anyone wishing to define a notion of assertibilty for an arbitrary sentence of the form "Square *s* is red." When is such a sentence assertible? When at least one person is willing to assert it, even though most others would deny it? When almost everyone is willing to assert it? Or when, say, two thirds are willing to assert it?

The reality is that people do differ in what they say and that we do not generally get into very serious messes as a result. Rather, since we do use these predicates, they must serve us somehow.

Mutual knowledge and a search problem

Lewis Carroll tells us *in Alice through the Looking Glass* that "T'was brillig and the slithy toves did gyre and gimble in the wabe; All mimsy were the borogroves and the mome raths outgrabe." I am certainly happy for the slithy toves who seem to have had a good time (or were they perhaps being tortured?), but what on earth *are* they and what *were* they doing?

Communication depends on shared language and this particular piece of language is not something we share with Carroll. But how does our usage of "red" or "bald" differ from that of "brillig"? It helps that we do all agree on *some* cases of things being red and of other things being not red and similarly with bald. If I have two books on the table such that both of us agree that one is red and that the other is not, then my

request, "Please give me the red book" will be satisfactorily met by you. The fact that we disagree on the colors of other books is not germane here.

Suppose however that there are three books, one for which both of us agree that it is red, one for which both of us agree that it is not red, and one which you, but not I are inclined to call red. If I ask you then to give me the red book, you might say "But there are two red books. Which one do you mean?", and I will explain which one, perhaps with the words "But the other one is not really red, but pink, wouldn't you say?"

In [VU] I show how a description of a book as "blue" can help another person to look for it, even though the speaker and the listener do not necessarily agree for every book whether it is blue or not. In the example in question, Ann asks her husband Bob to look for a certain book among a thousand on her shelf, after explaining to him that it is blue. It so happens that among these thousand books, there are 225 books that both would call blue; 25 that Ann, but not Bob, would call blue; 75 that Bob, but not Ann, would call blue; and 675 that neither would. Thus from Ann's point of view there are 250 blue books and from Bob's there are 300. When Ann says, "It is blue", meaning blue in *her* sense, Bob innocently assumes that it is blue in *his* sense. In spite of this fact, this dialogue between the two reduces Bob's expected searching time by 60%. It decreases from an (on average) search through 500 books (half of 1,000) without Ann's statement, to an (on average) search through 200 books.

The calculation goes as follows: if Ann had said nothing to Bob, he would need to search through *all* books. He might be lucky and the book that Ann wanted might be the very first. Or, in the worst case, it could be the very last book. So on average he would need to look through roughly 500 books.[13]

However, *with* Ann's statement that the book is blue, he will look first through all the books that *he* thinks are blue. There are 300 such books and of the 250 books that Ann considers blue, 225 are among these 300. So with a 90% chance, Ann's book, which she considers to be blue, will also be considered blue by Bob. In that lucky case, he will find it among these 300, and will only need to look through 150 books on average.

However, there is a 10% chance that the book, though blue for Ann is not blue for Bob, and he will not find it among the 300. In that case, after looking through the first 300, he will have to look at the other 700 books on Ann's shelves, and, on average, this second search will involve looking at 350 additional books. So with a 10% chance, Bob will look at 300+350 = 650 books. In sum, with a 90% chance he will search through 150 books, and with a 10% chance, through 650 books. Thus, his average comes out to.9 × 150 + .1 × 650 = 200 books. This is a very substantial reduction from 500.

In reality, Ann's information will be of more use to him. Even if he finds that Ann's book is not among the 300 that he thinks of as blue, he is unlikely to look at books that look clearly red or yellow to him. Most likely he will look only at books that *could* be called blue, and so in fact he will save more time than the 60% reduction that the calculation above shows.

One could reasonably say here *both* that communication took place *and* that Bob misunderstood Ann. Communication took place in the sense that Bob was helped by Ann's statement that the book was blue. But we can also say that Bob *misunderstood*

Ann in that he understood her to have intended that the book was blue in *his* sense, when in fact Ann had intended that it was blue in *hers*, and the two notions of blue were not the same.

The calculation, and the helpfulness to Bob of Ann's statement, do not depend on Ann and Bob knowing precisely how they use the word "blue", nor on such usage being the same from one occasion to another. As long as Ann's usage and Bob's usage are close enough, they may both differ from occasion to occasion and they will still be able to communicate in a useful way.

Let us make a calculation for the general case where the numbers might well be different. In the following, if X is a set of books with m elements, then $m/1,000$ is the probability $p(X)$ of X. $p(X \mid Y)$ is the probability of X given Y.

Let B_a, B_b be Ann's and Bob's versions, respectively, of the predicate "blue" and let N be the total number of Ann's books. Then (ignoring the .5) without any information from Ann, Bob would search through N/2 books. *With* the information from Ann, Bob will search through Nx books, where x is $(1/2)p(B_b|B_a)p(B_b) + p(\neg B_b|B_a)[p(B_b) + (1/2)p(\neg B_b)]$. Nx will be less than $N/2$ precisely when x is less than 1/2. We can show that the condition for Bob to profit from Ann's statement is that $p(B_b \mid B_a) > p(B_b)$,[14] or, in other words, that B_b and B_a have a positive correlation. This is a very lax requirement, quite far from requiring that $B_b = B_a$ and explains very well why we can find color words useful in communicating. Of course, usually $p(B_b|B_a)$ will be fairly close to 1, especially if Ann and Bob have been married for a long time, and the saving in time to Bob will be substantial.

Other contexts

Our use of vague predicates governs our actions and we may choose among actions according to whether we think that the predicates hold or do not. It is important then, that when a misunderstanding is possible, we avoid actions where the harm done by an incorrect action would be great.

Traffic lights are a good example of a situation where much rides on our using color words in the same way. When the light is green for me, it is red for you, and you must stop and I may go. An accident will be avoided only if we both agree which light is green and that the other light is red. So it is crucial here that we use colors which are not easily mistaken for one another. If traffic lights were instead scarlet, pink and crimson, meaning 'go', 'slow' and 'stop', respectively, it is easy to see that the number of accidents would be very much higher. It would be far more likely that two drivers on two different cross streets, disagreeing on the colors of the lights they saw, would both stop or both go forward. Our choice of red, yellow, and green for the three traffic signals, rather than scarlet, pink, and crimson indicates that we are somehow aware of the fact that our agreement on the differences among these last three is much more shaky than on the differences among the first three.

It is worth pointing out that the semantics of "scarlet" is no more, or no less, problematic than the semantics of "red" or "green". However, there is a clear difference in *application* between the triple (green,yellow,red) which we do use, and the triple (scarlet,pink,crimson) which we, very wisely, avoid.

Thus when we have a situation where different interpretations have a potential for disaster, we no longer rely on applying the vague predicates to *arbitrary* objects. Rather we confine the domain of application to a small set, like the set of the three lights, where disagreements are almost certain not to appear, or at least not to appear among drivers with normal vision.[15]

The linguistic practices of a society can be thought of as an enormous game, a game which is largely co-operative, as Grice[16] points out. Most of the time, when making statements to others, we seek to be informative, relevant, truthful, and brief. However, in this social game, each of us is playing a strategy which of necessity depends on our *own* interpretations of certain predicates, interpretations which may on occasion differ from those of other people and even from our own at other times.

Part of our skill in our usage of vague predicates is that we confine ourselves to strategies where differences in interpretation cannot lead to disasters and are more likely to be helpful.

Imagine, by contrast, a group of would-be bank robbers who agree to go into a bank, and agree also that when any old woman wearing a red dress enters the bank, they will all pull out their weapons and shout, "This is a hold-up." Confusion is likely and the probability of plans going awry is great if a middle aged woman wearing a magenta dress enters the bank. The robbers will not have anticipated possible disagreements on "red" and "old" -they will probably not have discussed whether magenta is a kind of red, or indeed what magenta looks like, and what ages are covered by "old". Thus it will probably happen that some of them will pull out their weapons while others are still dreaming about how they will spend their loot.

Actual robbers are unlikely to come up with a plan as stupid as the one I have described, but vague predicates applied to *non-cooperative* situations do sometimes lead to disaster. E.g., if two armies are meeting on a battlefield, then the weaker army should withdraw rather than risk defeat, regardless of the justice of its cause. However, since perceptions of "weaker" tend to differ, it is quite possible for each general, even if he is in full possession of the facts, to believe his own army to be stronger. Since both generals cannot be right, one of them will necessarily lead his men to defeat and to possible death.

Obviously it would be impossible for us to draw a precise line betwen those cases where agreement is great enough and those where it is not. While we often have an idea of when there is good agreement, we do not have precise data on how much agreement there actually *is*, nor a precise definition of how much is *needed*.[17] As Crispin Wright[18] says,

> It is (...) not merely that no sharp distinction may be drawn between cases where such a predicate applies and cases where it does not, but that no such distinction may be drawn between cases where it is definitely correct to apply the predicate and cases of *any* other sort.

Nash Equilibria

Let us suppose that three people -Alice, Bob, and Charles- are thinking of going to a restaurant to eat and they can choose between the Taj and the Peking, both of which are nearby and reasonably priced.[19] However, no one wants to eat alone and so if both Ann

and Bob are going to the Taj, Charles is best off also going to the Taj. Similarly, if both Bob and Charles are going to the Peking, Ann should also go to the Peking. Thus given their preferences, (Taj,Taj,Taj), where Ann, Bob and Charles respectively, all choose the Taj, is what is called a *Nash equilibrium* by game theorists[20]. (Peking,Peking,Peking) is the other Nash equilibrium in this situation. If two of the three diners are going to the Peking, the third should also go to the Peking.

But how would the three *establish* one of these equilibria? Presumably through a discussion which will no doubt involve the use of the words "Peking" and "Taj". Thus the words "Peking" and "Taj" will serve to help create the Nash equilibrium which, one should recall, is not an equilibrium of *words* but of *actions*.

Suppose that they have their discussion at 2 PM, agree to have dinner at the Peking, and then separate for the afternoon. There will then be a problem if Bob is under the impression that the Peking is the restaurant at Fifth Avenue and 43rd Street, which is actually the location of the Taj. In that case, the agreement "Let us all go to the Peking" will not result in actual satisfaction for all three of them. Ann and Charles will actually be at the Peking, waiting for Bob, while Bob will be either eating alone at the Taj or standing at Fifth and 43rd, wondering where the Peking is.

To ensure a satisfying outcome we need to insure not only that the three agree on the same *name* for the restaurant where they are eating, but also that, having agreed, they *apply* the name in the same way. Thus, what one might call a verbal Nash equilibrium, where they all agree to go to the Peking, will lead to a satisfactory outcome only if the three interpret the words in the same way.

As we have seen, if "Peking" and "Taj" had been vague words, then a misunderstanding would not only be possible, it would be almost inevitable. In such a situation, either one would need to take precautions to prevent an undesired outcome or else accept a less desired outcome if it is not too serious.

The example with the traffic lights is of the kind where we take precautions and make sure that the words are applied in the same way. If two cars driven by Ann and Bob are at an intersection, then the two Nash equiibria are (go,stop) and (stop,go), the first being one where Ann goes and Bob stops and the other being one where Bob goes and Ann stops. Now the first Nash equilibrium is created by the lights being, respectively, green for Ann and red for Bob. For a disaster to be avoided, it is important that (green,red) and (red,green) are the only two perceived situations, the first color being the one seen by Ann and the other being the one seen by Bob. We need the reliability of color perception, *in this limited context* in order to avoid the deadlock (stop,stop) and the disastrous (go,go).

Given the sharp difference between the colors red and green, it is unlikely that both Bob and Ann could believe that they are seeing a green light. However, it is quite likely that if the lights were "scarlet," "pink," and "crimson," then both of them might believe that they are seeing a scarlet light when in fact one of them was supposed to be seeing a crimson light and to stop.

For a different kind of example, suppose that Ann, wanting to go to a movie, asks Bob to suggest a romantic one. No doubt they will both agree that a Bruce Lee movie or a porno film is not what Ann has in mind. But what about *Basic Instinct*? Perhaps a discussion will reveal whether this is the sort of movie that Ann considers as falling

within her notion of "romantic." But if it does turn out that she goes to see *Basic Instinct*, thinking it to be romantic, and finds that it is not, that is not a disaster in the same way that a collision between two cars would be.

Why does truth matter to us?

Yet, just as the notion of truth has its opposite notion 'not true', so the notion of assertibility also has the complementary notion of non-assertibility. If "assertible" does not denote a predicate, and I have indicated earlier why it cannot, "non-assertible" must have some other use than as the mere negation of assertible. Let us see what such a use might be.

Suppose that a murder has taken place in a certain house, and a witness on the trial stand describes a man entering that house on the day of the murder as "bald." Independent evidence reveals that only one man entered that house on that occasion, that he still had 20% of his original hair, and that two-thirds of the community do not regard a man with 20% of his hair as bald. Can the witness be convicted of perjury?

Even if the sentence, "That man was bald" is not one that most of us *would* utter in such circumstances, it does not therefore follow that the witness is guilty of perjury. What she has said is not "beyond the pale." If on the other hand, the witness had described Mick Jagger as bald, she probably *would* be guilty of perjury. This is because, on hearing that the man was bald, we are likely to rule out Mick Jagger but are not likely to rule out someone who had 20% of his hair.

The non-assertibility of a sentence depends on how *we* are likely to interpret it, after making allowances for individual differences, and on what the consequnces are of any misinterpretations. Finally, it depends also on the intentions of someone who may have deliberately given us reason to choose a wrong action, which in this case may consist of convicting an innocent person or acquitting a guilty one.

The considerations above make it clear that with vague predicates there will be cases where certain sentences are *assertible* in the sense that we might ourselves assert them and other cases of sentences which are *non-assertible* in the sense that we ourselves (and many others) would reproach someone who used them. But there will also be the intermediate kinds of sentences, where we might *allow* their use; and whether we allow a particular use or not may well depend on the circumstances of the occasion.[21]

Conclusion

We still owe the reader an explanation of why we find the condition acceptable that if two books are indistinguishable in color then they are both blue or neither is. The fact that two books *are* indistinguishable does not guarantee that Bob and Ann will therefore agree on what color they are. But on any one occasion when Ann thinks of one of the books as blue, she will probably also think of the other as blue, and similarly for Bob. This is the sense in which the condition (iiia) at the beginning of the paper is justified.

Suppose that Ann is a person who owns very many books, and a series of books could be constructed from her collection such that the first of these books looks blue to

both Ann and Bob, the last does not look blue to either of them, and any two successive ones are indistinguishable in color by either Ann or Bob. If Ann becomes aware of this fact, she may realize that there is no consistent way for her to use the word "blue" and she might very well hesitate to use the word "blue" at all. I hope that the arguments I have offered will re-assure her. Ann does not need to be concerned with any abstract notion of truth or even with consistency. It is enough for her to know that in all probability Bob will be helped by what she has said to him.

Notes

1 Affiliated with the Department of Computer Science, Brooklyn College and the Doctoral programs in Computer Science, Mathematics and Philosophy, City University Graduate Center.

2 'Truth and Assertibility', *Journal of Philosophy* vol. LXXIII, 1976, 137-149.

3 Even the notion of assertibility will only be defended in a rather weak sense, and it will be apparent that a stronger sense is just as problematic as the notion of truth in this context.

4 A similar paradox can be generated by taking a man with a full head of hair and pulling off his hairs one by one till none is left. The paradox will arise because he is not bald in the beginning, he *is* bald at the end, and yet, removal of just one hair cannot make a person bald.

5 'Why there are no people', *Midwest studies in philosophy* 5, 411-467

6 'On that which is not', *Synthese* 41, 155-173

7 E.g. 'Fuzzy Logic and Approximate Reasoning', *Synthese* 30, 1975, 407-428.

8 See for instance, Fine, K.: 1975, 'Vagueness, Truth and Logic', *Synthese* 30, 265-300.

9 'Vagueness and alternative logic', in *Realism and reason*, Cambridge University Press, 1983, pp. 271-286. In this paper Hilary Putnam discusses several proposals, including those of Unger and Wheeler and an earlier proposal of my own which I made in 'Existence and Feasibility in Arithmetic', *Jour. Symbolic Logic* 36, 1971, 494-508. My earlier proposal amounted to accepting *local* reasoning while eschewing longer arguments and showing how we can make this distinction precise. The proposal in the current paper is not in conflict with the one I made in 1971, but emphasizes use rather than reasoning.

10 There is a very valuable special issue of the *Southern Journal of Philosophy* vol. 33, supplement, 1994, which is devoted to vagueness. There are papers by Michael Tye, Terry Horgan, Diana Raffman, Crispin Wright, Roy Sorenson, Timothy Williamson, and Vann McGee and Brian McLaughlin, discussing various alternatives. None of the papers discuss a utilitarian approach akin to mine, though Terry Horgan and Diana Raffman seem keenly aware of the occasion dependence of the usage of vague words.

11 'Vagueness and Utility: the Semantics of Common Nouns', *Linguistics and Philosophy* 17, 1994, 521-35. [VU] from now on.

12 I performed this experiment at different places, e.g. at CUNY, at Indiana University, and in Madras. The results were always similar. Diana Raffman in [VWP] 'Vagueness without paradox', *The Philosophical Review* 103, 1994, 41-74, also reports similar facts in section 5 of her paper.

13 Actually it is 500.5, the average of 1 and 1,000, but I shall ignore this .5.

14 $x = (1/2)p(B_b|B_a)p(B_b) + p(\neg B_b|B_a)[p(B_b) + (1/2)p(\neg B_b)]$

$= (1/2)[p(B_b|B_a)p(B_b) + p(\neg B_b|B_a)p(B_b)] + (1/2)p(\neg B_b|B_a)[p(B_b) + p(\neg B_b)]$

$= (1/2)p(B_b) + (1/2)[p(\neg B_b|B_a) = (1/2)[p(B_b) + [p(\neg B_b|B_a)].$

This last expression for x shows that x is less than 1/2 = (1/2)[$p(B_b) + p(\neg B_b)$] just in case $p(\neg B_b|B_a) < p(\neg B_b)$, or equivalently, $p(B_b|B_a) > p(B_b)$

15 Some color blind friends have told me that they tend to rely on the physical position of the three colors, which is apparently consistent.

16 Studies in the way of words, Harvard U. Press, 1989.

17 Indeed the experiment with the colored squares comes as a surprise to most people. There is hardly any awareness that there are individual differences in the use of color words, though there may be more awareness of individual differences with other words denoting vague predicates.

18 'On the Coherence of Vague Predicates', Synthese 30, 1975, 325-366.

19 The restaurants are fictitious.

20 Given three choices (u,v,w), they form a Nash equilibrium if for any other choice x, (x,v,w) does not have greater value for Ann than (u,v,w); for any y, (u,y,w), has no greater value for Bob than (u,v,w); and for any z, (u,v,z) has no greater value for Charles than (u,v,w). It is an equilibrium, since if two of the three hold on to their strategy (choice), the third has no incentive to shift. For a discussion of the Nash Equilibrium, see the article with the same title, by David Kreps in the The new Palgrave, a dictionary of economics, edited by Eatwell, Milgate and Newman, MacMillan, 1987. David Lewis' Convention, a philosophical study, Harvard University Press, 1969, contains a discussion of conventions which is very game theoretic in spirit and relevant to the current discussion.

11 Diana Raffman [VWP] also makes a strong case for the context dependence of vague predicates, though without our emphasis on utility.

Rohit Parikh is Distinguished Professor of Computer Science, Mathematics and Philosophy, Brooklyn College and City University Graduate Center. He has also taught at Stanford University, Panjab University, Bristol University and Boston University. Member of the Association for Computing and Association for Symbolic Logic. Former Chief Editor for the International Journal for Foundations of Computer Science. Program Chair for the Conferences *Logic in Computer Science* (1989), *Logic of Programs* (1985), and *Theoretical Aspects of Reasoning about Knowledge* (1990). Director of the Boston Logic Colloquium (1972-1981). Research publications in Recursion Theory, Proof Theory, Formal Languages, Non-standard Analysis, Dynamic Logic, Logic of Knowledge and Philosophy of Language.