

Unsharpenable Vagueness

John Collins and Achille C. Varzi
Columbia University

A plausible thought about vagueness is that it involves a form of semantic incompleteness. To say that a predicate is vague is to say (at the very least) that its extension is incompletely specified. And where there is incomplete specification of extension, there is indeterminacy—an indeterminacy between various ways that the extension of the predicate might be completed or, as some like to say, *sharpened* (precisified). We shall argue that this idea is defective insofar as there are vague predicates that cannot be sharpened. At least, there are predicates that are vague but that cannot be sharpened in such a way as to meet certain basic constraints that we think must be imposed on the very notion of sharpening.

I. TAKE IT OR LEAVE IT?

Consider the following version of a game known as *Take-It-or-Leave-It*, sometimes also referred to as the *Centipede*. There are two players, X and Y, and a game leader, or banker, who acts as a generous source of money, but otherwise takes no part in the proceedings. At the beginning of the game, the banker places \$2 on the table. Player X has the choice of taking the money or leaving it. If she takes it, the game finishes. If she leaves it, the banker adds another \$1 to the pot, and it is Y's turn to move. Y now has the

choice of taking the \$3 or leaving it. If Y takes the money, the banker compensates X by paying her \$1, and the game terminates. If Y leaves the money, then a further \$1 is placed on the table, and the choice reverts to X. If X takes the money this time, Y is compensated by receiving \$2 from the banker. Play continues in this fashion until one of the players takes the money. The other player is then compensated by receiving \$2 less than the amount taken. If both players continue to leave the money, play continues until there are \$101 on the table. If Y leaves the money at that point, it goes to X; Y gets his \$99 compensation instead. The game is illustrated in the following diagram. Nodes are marked 'X' or 'Y' according to whose turn it is to choose. Outcomes are indicated at the end of each branch, with X's payoff given first, and Y's second. The payoffs at the n th node always add up to $2n$. If n is even then the payoffs are $\langle n-1, n+1 \rangle$. If n is odd, the payoffs are $\langle n+1, n-1 \rangle$.




FIG. 1

Now consider the class of games which are final segments of the Take-It-or-Leave-It game just described. Let TL_n be the game that commences at the n th node of the diagram (our original game—the full Centipede—is TL_1). The game TL_{100} is completely trivial. It consists of a single node at which player Y gets to choose whether to take \$101 or leave it to X and receive \$99 compensation instead. Assuming that all our players care about is maximizing financial gain, Y will certainly take the \$101. But this suggests that the game TL_{99} is trivial as well. Here player X has the choice of pocketing \$100 or leaving it, and allowing Y to play the game TL_{100} . But allowing Y to play TL_{100} will certainly result in X receiving only \$99, so X is better off taking the \$100, and if she is rational she will do so. This in turn gives Y a compelling reason to take the money at his first move in game TL_{98} . This line of reasoning may be continued; it is usually referred to as the *backward induction argument*. If the rationality of both players is a matter of common knowledge, backward induction implies that the player whose turn it is to move first should take the money every time, whether the game is TL_{100} , or TL_{99} , . . . or TL_1 .

But this conclusion seems highly counterintuitive. Consider TL_1 , the original Centipede. According to the backward induction argument, player

X should terminate the game at the first move by taking the \$2. Y would then get nothing. But X realizes that she (and Y) would do much better if the game were to continue until there were, say, \$70 or \$80 on the table. At some point, of course, one of the players will cut and run. But if, for example, Y were to take the money when there were \$78 in the pot, X would still get \$76 in compensation—a far more attractive outcome than the mere \$2 guaranteed by backward induction.

II. RATIONALITY AND VAGUENESS

We take the above story to imply that rationality predicates are, to some degree, vague. This is so because the backward induction argument can be reconstrued as a sorites argument. Call a TL game a *take-it game* if the first player to move is rationally obliged to take the money. Similarly, call such a game a *leave-it game* if considerations of ideal rationality permit the first player to leave the money on the table. TL100 is obviously a take-it game. And we think it is just as obvious that TL1 is a leave-it game. However, there appears to be no clear boundary between take-it and leave-it games. One dollar less on the table hardly makes a difference to a player's preferences. Thus we seem forced to reason as follows:

TL100 is a take-it game.
If TL($n+1$) is a take-it game, then so is TL n .
Ergo, TL1 is a take-it game.

And this is a sorites argument leading in one hundred steps from plausible premises to a false conclusion. It shows that the predicate 'take-it game' is vague, just as the argument

A 100 year-old person is elderly.
If an $n+1$ year-old person is elderly, then so is an n year-old person.
Ergo, a 1 year-old person is elderly.

shows that the predicate 'elderly' is vague. A similar argument can of course be constructed to show that the predicate 'leave-it game' is also vague. And since any term involved in the definitions of 'take-it game' and 'leave-it game' is (or can be assumed to be) perfectly precise, except for the predicates 'rationally obliged' and 'rationally permitted', respectively, this means that the burden of the sorites lies entirely on the latter predicates.

It is important to realize the role of the inductive premise in the soritical chains generated by the Centipede. One can easily think of sequences of rationality claims involving marginal change, but generally considerations of optimization can determine the relevant cut-off point. For instance, it is true that one is rationally obliged to put at least 1 percent of one's monthly

salary into a retirement plan, and it is false that one is rationally obliged to put at least 100 percent of one's monthly salary into a retirement plan. Yet, any sequence obtained by filling in a suitable number of intermediate steps is likely to violate the relevant soritical induction. The conditional

If one is rationally obliged to put at least k percent of one's monthly salary into a retirement plan, then one is rationally obliged to put at least $k+1$ percent of one's monthly salary into a retirement plan.

may seem intuitively true regardless of the specific value of k , but closer examination will reveal the value of k beyond which there is no rational obligation to sacrifice one's salary in favor of one's retirement plan. Not so in the case of our Centipede. In such case, no considerations of optimization would seem to be of any help. The Centipede yields a true sorites.

III. UNSHARPENABLE VAGUENESS

That rationality predicates are vague is bad news, but it may come as no surprise. Theories of rational choice hardly aim at the completeness that seems required to protect them from this outcome. Some may be prepared to go the tough way, of course. In the case of the Centipede, some may be prepared to bite the bullet and accept the conclusion that the only rational first move for a player is to take the \$2. But this sounds to us like the response of the person who, led through a sequence of ten thousand colored tiles each visually indistinguishable from its predecessor, finds herself at the end of the sequence staring at a tile that is manifestly red and maintaining that it looks yellow. Thus does the grip of theory strangle common sense.

To repeat, then: the predicates 'take-it game' and 'leave-it game' are vague. There is, however, something peculiar about the vagueness exhibited by these predicates, i.e., in the end, by the predicates 'rationally obliged' and 'rationally permitted'. For this vagueness appears to be *unsharpenable*. As we shall argue below, the extensions of these predicates appear to have vague boundaries which resist any attempt at making them sharp, if only for the purpose of considering the available options. And this is peculiar, because sharpenability is generally regarded as a natural feature of vague predicates. To say that a predicate is vague is to imply that its extension is incompletely specified, hence that there is a certain indeterminacy between various ways that the extension of the predicate might be completed or *sharpened*.

Various theories of vagueness rely heavily on this feature of vague predicates. The view that vagueness is merely a sign of semantic laziness, for instance, rests on the idea that the indeterminacy can always be elimi-

nated, at least in principle. And supervaluationism, for another example, exploits the idea that the truth value of a sentence involving vague predicates is ultimately a function of the classical truth values that the sentence receives under the various ways in which those predicates can conceivably be sharpened. If the sentence receives the same value, say True, in *every* such sharpening, then—so goes the theory—the vagueness of the predicates is ultimately irrelevant, and the sentence should be assigned that value. If, on the other hand, the value turns out to vary from sharpening to sharpening, then the vagueness appears to be irredeemable and the sentence cannot receive a definite truth value.

Both the vagueness-as-laziness theorist and the supervaluationist are, of course, aware that in some cases it may be utterly difficult, if not impossible, to deal with actual sharpenings. For instance, it may be utterly difficult, if not impossible, to measure someone's height with enough accuracy to establish whether or not she qualifies as tall relative to a certain admissible sharpening of the vague predicate 'tall'. But this is tolerable. As long as the sharpening guarantees the existence of a fact of the matter, logic is saved and the only difficulties that may be left concern epistemology. After all, that is one reason why the language we speak is not precise: making it precise would be very costly, and yet impractical.

The case of the predicates 'take-it game' and 'leave-it game' is different, though. When we say that the vagueness exhibited by these predicates is unsharpenable, we mean more than a pragmatic or epistemic difficulty. We mean to say that with these predicates it is *in principle* impossible to come up with any admissible sharpening at all. These predicates have vague boundaries that *as a matter of principle* resist any attempt of making them sharp. And this may not be tolerable for a theory of vagueness that relies on sharpenability.

IV. WHY THESE PREDICATES ARE UNSHARPENABLE

Let's see why that is so. In order for a sharpening to be admissible, various kinds of condition must be met. One sort of condition involves all those constraints of the kind that are known as penumbral connections. For example, any sharpening of the predicates 'pink' and 'red' should respect the constraint of mutual exclusiveness: nothing can be both pink and red. So a sharpening of these predicates should never yield overlapping extensions. Likewise, any sharpening of 'red' should respect the internal constraint that if a borderline object x is classified as red, and another borderline object y has a shade that falls between that of x and that of a clear red object, then y should be classified as red as well.

A second sort of condition—we submit—is a requirement of *public accessibility*. By this we mean the following. Suppose that *V* is a vague predicate and *V** is a potential sharpening of *V* that meets every constraint of the first kind. Still *V** will only be an admissible sharpening of *V* if it is in principle possible for any two speakers of the language to shift their standards of correctness so as to accord with the rule for proper application of *V**. It must in principle be possible, in other words, for our two speakers to decide to speak the language in which *V** replaces *V*, and for it to be common knowledge that this shift has taken place. (Let us emphasize that this possibility must be available as a matter of principle. In practice, as we have seen, this may not be the case.)

An example will help illustrate this point. Consider the vague predicate ‘rich’. We might have some reason to precisify this term, at least as it appeared in a certain context. Suppose, for example, that some community of speakers of our language comes to agree that it would be a good idea for those who are rich to pay 10 percent more tax. If this is to be passed into law in the form of some official edict:

Henceforth those who are rich shall be taxed at a rate 10 percent higher than the current rate,

then we must either (i) replace the vague term ‘rich’ with more precise language or (ii) retain the word ‘rich’ but settle upon some way of sharpening its meaning. Let’s focus on the second method. Let us assume, for the sake of simplicity, that all the speakers in the community of our example are equally competent in their ability to use the predicate ‘rich’, and let us assume further that any sharpening of the term must be such as to leave all members of the community equally competent in their abilities to apply the newly revised replacement. Now suppose that, contrary to the public accessibility requirement, the method of sharpening in fact proceeds in the following way. Each of the speakers in our community individually refines the meaning of the vague term ‘rich’ in his or her own idiolect by sharpening its extension. And suppose that, as a matter of fact, no speaker knows the details of any one of these private sharpenings other than his or her own. (This is, of course, a much stronger claim than the simple denial of the common knowledge assumption, but it will be useful, and sufficient for our purposes, to consider how the argument would go in this extreme case.) This manifold revision of meaning has now, we shall imagine, taken place. What are we to say about the meaning of the predicate ‘rich’ as it is now used by the community?

There seem to be two possibilities. On the one hand, we might say that the original predicate has simply fragmented into a bewildering array of more or less closely related, homophonic, yet semantically distinct predicates with a range of different meanings. That is to say, after the individual

sharpenings are complete, you, she, and I simply mean different things when we use the word 'rich'. But if that is the case, then what we as a group have achieved certainly does not amount to a sharpening of the original vague predicate. We have, rather, merely succeeded in replacing vagueness by ambiguity. On the other hand, if we insist that after the individual sharpenings have been performed, all of us still mean the same thing when we use the word 'rich', then it is hard to resist the conclusion that the predicate remains vague in spite of our individual efforts at precisification. For various speakers will have privately sharpened the predicate in various different ways, and clearly the collection of all of these private sharpenings must be taken to demarcate an extended penumbral region for the revised predicate insofar as its new meaning is determined by the way the word is now used throughout the community. Otherwise we would have to admit that after the revision of meaning has taken place, some speakers have become more accurate in their use of the predicate than others, and this would contradict our assumption that all of the speakers in the group are, and remain, equally competent users of the predicate in question. Hence, the public predicate 'rich' is still vague. The argument we have just given rules out, for example, the possibility that a range of individual sharpenings might determine a single community-wide standard precisification via some kind of averaging method. *A family of precise private languages does not constitute a single precise public language.*

Now, the requirement of public accessibility may be too strong as a general condition to be imposed on the sharpenings of *all* predicates. Some may think that, for example, a predicate such as 'handsome' simply does not admit of publicly accessible sharpenings: X may be willing to count as handsome persons whom Y is not willing to count. In this respect, 'handsome' would not suffer from the same kind of vagueness as 'rich'. However, in the case of rationality predicates such as 'take-it game' or 'leave-it game', public accessibility is hardly negotiable. If the standards of rationality are not to be found in the eye of the beholder, then it must always be possible for any two speakers of the language to shift their standards so as to accord at least in principle with the rule prescribed by an admissible sharpening of 'take-it game' or 'leave-it game'. Admissibility here presupposes accessibility.

But that can never be the case. Here the requirement of public accessibility will inevitably collide with one of the requirements of penumbral connection. What requirements of penumbral connection apply to, say, the predicate 'take-it game'? One obvious condition is the converse of the inductive premise of the sorites series: if TL_n is a take-it game, then so is $TL_{(n+1)}$. But we can derive a further constraint besides this one. That is because the vagueness of 'take-it game' derives from the vagueness of 'rationally obliged' and inherits the logical properties of the latter expression.

Even if the term ‘rationally obliged’ is, as we are arguing here, a vague one, various things remain non-negotiable about our concept of rational obligation. For example, if you face a choice between two alternatives A and B , and you are certain that the result of your choosing A would be better for you (by your own lights) than the result of your choosing B , then the choice of A is, for you, rationally obligatory. In the present case, that means, for example, that if you are faced with making the first move in TL_n , and you are certain that you would end up with more money by taking the cash than by leaving it, then you are rationally obliged to take the money. That is to say, if you find yourself (under the usual conditions of common knowledge of rationality) in the position of making the first move in the game TL_n , and you are certain that you will end up richer by taking the money rather than leaving it, then TL_n is a clear case of a take-it game. This is a fact about the correct application of the term ‘take-it game’ that counts as penumbral connection.

Yet this constraint and the condition of public accessibility cannot be met simultaneously. Sharpening the term ‘take-it game’ amounts to determining what is to count as the smallest value of n for which one is rationally obliged to take the money on the table at the first move of TL_n . Call this the *critical value* of n . Suppose that X and Y are considering a sharpening of ‘take-it game’ to the effect that the critical value of n is set at 50. Then TL_{50} is a take-it game, but TL_{49} is not. But now, when confronted with TL_{49} , X has the choice of taking \$50 or leaving Y with the first move in TL_{50} . If it is common knowledge that TL_{50} is to be counted as a take-it game, then X knows that Y will take the \$51 in that situation, leaving X with \$49 compensation. This means that when faced with making the first move in TL_{49} , X is certain that she will end up richer if she takes the money than if she leaves it. But by the condition of penumbral connection derived above, this means that she is rationally obliged to take the money in TL_{49} , i.e. that TL_{49} is also a take-it game. This is a *reductio* of the assumption that it can be common knowledge that the critical value of n is 50. And clearly the same argument applies no matter what value of n we settle on. Player X is only willing to consider sharpenings where n is odd; player Y only sharpenings where n is even.

V. CONCLUDING REMARKS

This argument undermines the idea that vagueness goes hand in hand with sharpenability. And if we are not mistaken, this fact is bad news for any theory exploiting that idea for the purpose of providing an account of the semantics of a vague language (such as supervaluationism or the vagueness-

as-laziness view). Of course, at this point one looks for ways of resisting the argument. Two objections, in particular, are worth considering.

The first objection concerns our very starting point. We admit that there are clear cases of take-it games; clear cases, that is, in which the first player is rationally obliged to take the money at the first move. This, one might argue, is to presuppose that the players are only concerned with maximizing utility in the short-term, ignoring the advantages that might accrue in the long term through cooperating and building trust.

We reply that if there is indeed a distinction to be made here between a player's short-term and long-term interests, then that player is not really playing a TL game—hence the question of whether she is playing a take-it game or a leave-it game does not even arise. For, the predicates 'take-it game' and 'leave-it game' are only defined in terms of the rational obligations of an agent playing a TL game. And by definition, an agent is only playing a TL game if her utilities for the various outcomes are as given in the game tree. Granted, it may be a serious question which, if any, real-life interactions are adequately modeled by the Centipede. It may be fair to wonder, that is, whether the utilities of real-life players coincide with those specified for the TL games. But this is irrelevant to our argument.

The second objection concerns the possibility of higher-order vagueness. We have assumed that there are clear cases of take-it games. But our argument presupposes, in addition, that every TL game can be classified either as a clear case or as a borderline case. This is why we can speak of admissible sharpenings in the first place. However, this may be questioned. There may be intermediate cases where it simply is indeterminate whether the game is a take-it game or a borderline case of a take-it game. That is, the predicate 'take-it game' may have borderline borderline cases. (Likewise for 'leave-it game'.) This is a common feature of vague predicates: they do not determine a sharp partition into the positive instances and the negative instances, but neither do they determine a sharp partition into the positive instances, the negative instances, and the borderline cases. Two cut-off points are not easier to find than one. And if things are so, then the public-accessibility requirement seems too strong as stated. It may be legitimate to expect any two speakers of the same language (any two players of our game) to agree on what counts as an admissible sharpening with regard to the *clear* borderline cases. But if different speakers (players), or even the same speaker (player) in different contexts, may have conflicting views as to whether a certain item qualifies as a borderline case, then how can they be expected to agree on what qualifies as an admissible sharpening of *the* borderline cases?

We do not think our argument is affected by this sort of concern either. Certainly the idea of a complete sharpening is a simplification, but the

argument would go through even if the notion of a sharpening turned out to be vague (or vaguely vague, and so on). Only, if that were the case, the public-accessibility constraint would have to be weakened so as to hold only for the paradigmatic cases of sharpening. Not every sharpening, that is, but at least every clear case of sharpening would have to be such as to qualify as admissible only if any two speakers of the language would agree on considering the relevant standards of correctness. This is quite reasonable even in the presence of higher-order vagueness. For instance, the fact that 'tall' is higher-order vague does not prevent this predicate from admitting of various indisputable sharpenings (corresponding to cut-off points lying somewhere in the middle of the area corresponding to the clear borderline cases). But our argument shows that the predicates 'take-it game' and 'leave-it game' do not admit of *any* such way of shifting the standards. So either these predicates involve no clear borderline cases at all, or they admit of some borderline cases and yet allow for no clear sharpenings. Either way, something has gone wrong with the idea that a semantics of vagueness must exploit the sharpenability of vague predicates. And that was the point of our argument.

NOTE

An early version of this paper was read at the *Conference on Methods in Philosophy and the Sciences* in New York, December 6, 1997. We are thankful to Stephen Schiffer for his comments and to the audience for a spirited discussion. Thanks also to Kit Fine, Haim Gaifman, Patrick Grim, Isaac Levi, Michael Morreau, Graham Priest, Roy Sorensen, and Jason Stanley for lively exchanges on these matters.