VAGUENESS, IGNORANCE, AND EPISTEMIC POSSIBILITIES

Zoltán VECSEY

ABSTRACT: The paper focuses on a hitherto unexamined version of the third possibility conception of vagueness. It is claimed that statements about borderline cases can be treated by analogy with statements about epistemic possibilities. The proposed account can be readily subsumed under the generic category 'third possibility view' because, in contrast to definitively true and definitively false application cases of vague predicates, statements about borderline cases are interpreted as non-truth-functional.

KEYWORDS: vagueness, ignorance, epistemicism, third possibility views

I. Introduction

In order to explain away the apparently paradoxical features of borderline cases, third possibility theories of vagueness are typically forced to introduce non-classical truth values into their semantics. Theories that make use of the Strong Kleene evaluation scheme or the method of supervaluations have to solve the puzzles of borderline statements by postulating truth value gaps. According to these theories, statements about borderline cases of a vague predicate come out true on some valuations and false on others, and are thus neither definitely true nor definitely false. One reasonable way to recover this semantic deficit is to conceive gappy statements as representing a third truth value, say ½. Borderline statements are then assigned the value ½.

Something similar happens in the case of paraconsistent accounts of vagueness. Theories that are committed to standard systems of paraconsistent logic have to admit truth value gluts in their formal semantic frameworks. Dialetheists, for example, permit borderline statements to be both true and false. This form of semantic anomaly is thought to be effectively resolved by the introduction of a non-classical truth value 2. While classical values 1 and 0 represent definitely true and definitely false statements, respectively, the third value 2 has the function of representing borderline statements that are supposed to fall in the intersection of definitive truth and definitive falsity.

Zoltán Vecsey

Theories of vagueness based on many valued logic differ from the accounts mentioned above in that they operate with a set of non-classical truth values instead of only one. In many valued settings borderline statements can take any member of the set of real numbers in the closed interval [0, 1] as they truth value. Such numerical values are often equated with degrees of truth. The main idea is that the higher numerical value a borderline statement has, the closer it is to definite truth, and similarly with lower values and definite falsity.

On my view, neither of the presently available third possibility solutions for dealing with the problem of borderline cases is entirely satisfactory. There is a general argument against them that goes something like this.[1] Predicate vagueness seems to indicate the presence at least of three things. First, if F is vague, then there must be cases where the application of F is definitely true and cases where the application of F is definitely false. Second, if a sorites series is created for F, then there must be a seamless transition between the cases of definitely true and definitively false applications. Third, there must be a borderline area of cases where the applications of F are neither definitely true nor definitively false. The question arising from this quick characterization is the following: Is there a noncontradictory way to find a proper semantic classification for borderline cases between the poles of definite truth and definite falsity? In positing a novel type of truth value – ½, 2, or real numbers in the closed interval [0, 1] –, third possibility theories answer the question in the affirmative. A crucial problem with this kind of answer is, however, that it is not in line with the seamlessness of the transition between contrasting cases of applications. More concretely, if borderline applications were assigned a third type of truth value that is strictly incompatible with definite truth and definite falsity, seamless transitions would have to be regarded as impossible. Instead of seamlessness, one would be confronted with sharp demarcations between true or false applications and applications that are gappy, glutty, or have an intermediate degree of truth. That would imply that we should categorically deny the existence of one of the most basic phenomena of predicate vagueness.

Given the simplicity and persuasiveness of this argument, I am inclined to think that all existing third possibility theories of vagueness should be abandoned.

---

[1] For further details see Crispin Wright, "On Being in a Quandary: Relativism, Vagueness, Logical Revisionism," *Mind* 110 (2001): 45–98, and "The Illusion of Higher-Order Vagueness," in *Cuts and Clouds: Vagueness, Its Truth, Its Nature, and Its Logic*, eds. Richard Dietz and Sebastiano Moruzzi (Oxford: Oxford University Press, 2010), 523-549, Stephen Schiffer, *The Things We Mean* (Oxford: Clarendon Press, 2003), and Matti Eklund, "Vagueness and Second-Level Indeterminacy," in *Cuts and Clouds*, 63-76.

At the same time, I do not think that the argument suffices to show that any such theory is untenable. As it frequently happens in other branches of fundamental linguistic-philosophical research, the theoretical possibilities in this domain of investigation are not yet fully explored. The remainder of this paper will focus on a hitherto unexamined version of the third possibility conception of vagueness. It will be claimed that statements about borderline cases can be treated by analogy with statements about epistemic possibilities. The proposed account can be readily subsumed under the generic category 'third possibility view' because, in contrast to definitively true and definitively false cases of application, statements about borderline cases will be interpreted as non-truth-functional.

II. "It is definitively the case that *Fa*"

Most contributors to the vagueness debate are of the opinion that the definitely operator ought to play a central role in the characterization of borderlineness.[2] It is not too surprising, however, that there is little agreement about how to specify that role in a generally acceptable way. One of the reasons behind the disagreement is that the views about the relationship between definiteness and truth differ significantly as we move between competing sides of the debate. Those who hold that borderline applications of vague predicates must be associated with some kind of non-classical truth value like to try to persuade us that the definitely operator should be interpreted in semantic terms. Consider the example of supervaluationism. Supervaluationists typically hold that the application of F to a is definitely true only if Fa comes out true under all semantically admissible evaluations. On the other hand, they suggest that definiteness and truth come apart when a counts as a borderline case of F according to some evaluations. In these cases, they contend, definitely Fa has to be evaluated as false, while Fa has to be assigned the value ½. Note that the the definitely operator displays here a splitting behavior. In clear cases of application it appears to be inextricably linked to the truth of the predicate it modifies. Suppose poor Fred has zero hairs on his head. Then the statement 'Fred is bald' comes out true according to all semantically admissible systems of evaluation, and so is definitely true. Beyond the undisputably clear cases, however, where the application of Fa becomes neither true nor false, definitely Fa should be regarded as false. Take now the case of borderline bald Felix who has 1025 hairs on his

---

[2] One notable exception is Rosanna Keefe, who argues that the (technical) question of definiteness does not belong to the central part of the theory of vagueness. See her *Theories of Vagueness* (Cambridge: Cambridge University Press, 2000).

head. According to supervaluationists, the statement "Felix is bald" is gappy, but the statement "It is definitively the case that Felix is bald" is false. This is obviously implausible. The definitively operator is similar in one sense to the generalized quantifiers 'some' and 'all.' These are technical terms that have well-understood counterparts in most natural languages. And if the semantics of these counterparts dictates uniform behavior in relevantly similar contexts, then it is also reasonable to require that the technical terms should behave uniformly in relevantly similar contexts. The contexts of our natural language statements about Fred and Felix are similar in the sense that they presuppose a relatively strong correlation between the properties of being definitely bald and being bald.[3] If being definitely bald as applied to Fred or Felix gets assigned a polar truth value in a given context, then it is natural to expect that being bald also gets assigned a polar truth value in the same context. As we have seen, supervaluations directly contravene that rule. Hence, at least intuitively, we may conclude that the supervaluationist semantics for 'definitely' is incorrect in its present form.

The diagnosis given above generalizes across all versions of non-classical treatments of borderline cases. Friends of paraconsistent logic and degree theorists are surely not in a much better position with respect to the clarification of the role of the technical term 'definitely.' These approaches have a common core in that they use the definitely operator in order to demonstrate that borderline applications of vague predicates must be associated with some kind of non-classical truth value. But this does not follow immediately from the intuitive meaning of 'definitely.' Nor does it follow that borderline application cases must give rise to definiteness in any sense. It would then seem better to have an interpretation of the definitely operator which does not involve third type truth values.

One attractive option in this regard is the epistemicist view worked out in details by Timothy Williamson and Patrick Greenough.[4] Epistemicism is well-known for its full preservation of classical logic and its bivalent semantics for

---

[3] It has to be noted that 'definitely' may be used either as a predicate modifier (*a* is definitely *F*) or as part of a sentence operator (it is definitely the case that *Fa*). Although it is not entirely self-evident, I will assume below that 'definitely' produces exactly the same semantic effects in both cases.

[4] See Timothy Williamson, *Vagueness* (London: Routledge, 1994), and "Reply to McGee and McLaughlin," *Linguistics and Philosophy* 27 (2004): 113–122, and Patrick Greenough, "Vagueness: A Minimal Theory," *Mind* 112 (2003): 235–281. For the critique of the episte-micist position see Stephen Schiffer, "The Epistemic Theory of Vagueness," *Philosophical Perspectives* 13 (1999): 481-503, and Zoltán Vecsey, "Epistemic Approaches to Vagueness," *Dialogue* 49, 2 (2010): 295–307.

vague discourse. Since the usual principles of classical reasoning leave no conceptual room for postulating third type truth values in the semantic machinery, epistemicists are in a position to provide a conservative-style explanation for the definitely operator. According to their view, the phenomenon of definiteness, as it appears in ordinary epistemic situations, can be exhaustively explained in terms of knowledge. The basic idea may be roughly stated as follows:[5]

> DEFINITENESS: a's being definitely F consists in the absence of obstacles to knowing that a is F.

Reflecting on this proposal, one may wonder whether "It is definitely the case that Fa" has the same epistemic status as "It is known that Fa." It depends. Williamson himself would argue against the identification of 'definitely' and 'knowably.' His ground for this is that in reasoning with vague predicates, a certain kind of epistemic uncertainty becomes inevitable. The source of the uncertainty is that we are not able to discriminate between cases of F-ness that are only marginally different. If two objects are so similar that we do not have any chance to distinguish them with respect to the instantiation of the property of F-ness, then we justly believe that both are F. But in cases where one of the indistinguishable objects is in fact F and the other is not-F, our beliefs are not reliable enough to count as knowledge. Williamson thus comes to the conclusion that some of the obstacles to knowing that a is F may prove ineliminable, even under optimal epistemic conditions. It seems, then, that we have to add a significant restriction to the explanation of the definitely operator:

> RESTRICTED DEFINITENESS: Since some of the obstacles to knowing are ineliminable, "It is definitely the case that Fa" cannot be epistemically equivalent to "It is known that Fa."

One question immediately arises: If 'definitely' does not collapse into 'knowably,' then why should we think that definiteness can be exhaustively, or at least adequately, explicated in terms of knowledge? Williamson's answer would be that RESTRICTED DEFINITENESS is deliberately vague, because definiteness itself is a vague phenomenon.[6] And this is why it would be folly to try to describe the relationship between definiteness and knowledge in a more rigorous or transparent manner. But now the same basic question arises again: If 'definitely' is

---

[5] Cf. Williamson , "Reply," 118.
[6] Williamson , "Reply," 118.

indeed a vague expression, then how can it be effectively applied in the epistemicist's analysis of borderline cases? In order to circumvent the difficulty implicit in this question, one may perhaps argue, following Greenough's minimal theory, that DEFINITENESS is superior to RESTRICTED DEFINITENESS, since it enables to express the epistemicist's central insight in a more theory-neutral way. Greenough contends, contra Williamson, that 'definitely' and 'knowably' may be taken to have the same meaning.[7] This is tantamount to acknowledging that truth is not entirely beyond our cognitive reach in undisputably clear application cases of vague predicates. Yet, interestingly enough, the theory Greenough advances does not sanction the acceptance of DEFINITENESS. Rather, it proposes to dispense with the definitely operator altogether. The supposed advantage of this move is that in this way it may become possible to elaborate a minimal conception of vagueness in purely epistemic terms. Perhaps Greenough is not completely wrong on this latter point. But in my view, the rejection of DEFINITENESS leads in the end to an incomplete and hence unsatisfactory theory of vagueness.

The reason why we should insist on the interchangeability of 'definitely' and 'knowably' is remarkably simple. When we say that "It is definitely the case that Fa," we are assuming that a instantiates a certain set of properties that are jointly necessary and sufficient for being F. It would be inconvenient to apply 'is definitely F' to a, if we were not entirely confident that there is no room for error about a's instantiation of F-ness. This may be regarded as the default epistemic assumption concerning our ordinary criteria for the application of 'definitely.' Consider again Fred, who has zero hairs on his head. Given that he is an adult male who has lost all of his hairs because of the natural process of aging, it would be quite pointless to debate that he is definitely bald. In situations like this, where a's instantiation of F-ness is beyond any reasonable doubt, the role of the definitely operator consists in ascribing epistemic necessity to Fa, and this, in turn, indicates that we are in a position to know that Fa is the case.[8] Therefore, if we are entirely confident that 'is definitely F' can be applied correctly to a, then we know that a is F.

One might object that the expressions 'entirely confident' and 'correctly' are obviously vague in the last sentence, so there is still no reason to defend the epistemic equivalence between "It is definitely the case that Fa" and "It is known that Fa." The objection can be easily answered by stating the argument in a less

---

[7] Greenough, "Vagueness," 251–252.

[8] Cf. Michael Huemer, "Epistemic Possibility," *Synthese* 156 (2007): 120.

informal way. Let us say that 'definitely' and 'knowably' are equivalent epistemically if and only if each of the following conditions is satisfied:

> i. Speakers of a community c are competent in using and understanding statements containing the predicate 'is definitely F.'
>
> ii. The predicate 'is definitely F is used in a transparent epistemic situation.[9]
>
> iii. There is no doubt on the part of the speakers of c that 'is definitely F applies to a.

Of course, one might continue to worry about the presence of implicit vagueness in the extensions of such terms as 'community,' 'competent' and 'understanding', etc. Moreover, one might complain that with the possible exception of the vocabulary of arithmetic every other expression, including 'vague,' is inherently vague.

The best reply to the first worry is to note that it is far from self-evident that the vagueness of 'community' and the likes is of the same semantic kind as the vagueness of 'bald.' The extension of 'community,' for example, seems to lack sharp boundaries, but it would be quite difficult to use it in a typical sorites argument. With respect to second complaint, it can be noted that conceiving vagueness as a pervasive phenomenon endemic to the vocabulary of most languages would reduce dramatically the prospects of coherent theorizing in this domain of research. If our semantic apparatus would indeed be thoroughly vague, then even such theoretical statements were infected with vagueness which are intended to express the pervasive vagueness of vague languages. That would be a bad consequence. I think this line of reply is persuasive enough to reject the above complaints as ill-motivated. Therefore, I take it for granted that there is no vagueness in the conditions i–iii.

So far it has been argued that in a broadly epistemicist framework undisputably clear application cases of vague predicates may be taken as implying knowledge. But nothing has been said about how this insight can help us in the task of characterizing borderlineness. The first step in this direction would be to draw an accurate distinction between cases where F definitely applies to a and cases where it is indeterminate or indefinite whether or not F applies to a. In

---

[9] Under 'transparent epistemic situation' I mean a situation which is not threatened by knowledge-scepticism and Gettier-free.

drawing this distinction, we may rely on those observations which we have already made on the equivalence between 'definitely' and 'knowably.' So we can say that in cases where it is indeterminate or indefinite whether or not *F* applies to *a*, conditions i and ii are satisfied, but condition iii is not satisfied. This means that speakers of a community *c* are competent in using and understanding the predicate 'is definitely *F*,' but they are uncertain whether or not it applies to *a* despite the fact that the epistemic situation they are in is sufficiently transparent. The presence of uncertainty concerning the applicability of 'is definitely *F* to *a* signals explicitly that *a* is not known to be *F* in *c*. On this ground, it can be argued that the difference between clear application cases of *F* and indeterminate or indefinite application cases of *F* depends ultimately on the presence or absence of knowledge on the speakers' part.

At this point, a potential misunderstanding needs to be avoided. It seems reasonable to assume that indeterminate or indefinite application cases of the predicate *F* may be interpreted as borderline cases of *F*-ness. It also seems reasonable to assume that in borderline cases speakers become uncertain of the applicability of 'is definitely *F* to *a* because they do not know whether or not *a* is *F*. In light of this, it is tempting to try to characterize borderline cases in terms of absence of knowledge.[10] But this is not the same as trying to provide a clear definition of borderlineness. Borderline cases cannot be defined on the ground of a clear definition of clear cases, since it would mean that there is a sharp tripartite division between *F*-ness, not-*F*-ness and a middle category in between. Such sharp divisions would be in conflict with the acknowledged seamlessness of the sorites transitions. In this respect, I am agreeing with Diana Raffman, who insists that there can be no definitely borderline cases of *F*-ness.[11] Despite this, we are able to recognize *a* as a borderline case, because we know that we are uncertain whether

---

[10] It must be added that many theories of vagueness refuse to equate borderline cases with absence of knowledge. The reasons behind the refusal remain, however, in most cases rather obscure. For example, adherents of the psychological theory of vagueness contend that borderline cases have to be explained in terms of the quandary state of speakers. According to this view, quandary is a state of ambivalence rather than ignorance. Unfortunately, *it is not quite clear why* the psychological state of ambivalence should be seen as different in kind from the mental state of ignorance. See Wright, "On Being in a Quandary," and Stephen Schiffer, "Vague Properties," in *Cuts and Clouds*, 109-130.

[11] Diana Raffman, "Demoting Higher-Order Vagueness," in *Cuts and Clouds*, 513. In this paper, Raffman argues that from the impossibility of definitely borderline cases of *F*-ness it follows that there are also no borderline borderline cases of *F*-ness. Raffman's argument, in my opinion, is effective against any theory which would go beyond first-order vagueness.

or not the predicate F applies to it. Being a borderline case of F-ness in this way presents itself as an elusive property that cannot be clearly defined in terms of positive or negative definite cases of F-ness.

The last remark, I must admit, reveals relatively little about how borderline statements should be evaluated in the present framework. If borderlineness is really an elusive property generated by speakers' ignorance, then it is not evident how can statements about borderline cases be assigned any classical semantic value. And if it turns out that these statements cannot be known to be true or false in the classical sense, then the question arises as to how to handle them without invoking non-classical truth values. These are the issues I will deal with in some detail in the next section.

III. Borderline statements and epistemic possibilities

Let us turn back to the example of Fred who has zero hairs on his head. Suppose conditions i–iii are satisfied. Then the statement "Fred is definitely bald" should be assigned the polar truth value true in community c. And given that the inference from 'definitely F' to 'F' is valid in standard epistemic logic, the statement 'Fred is bald' will be also true in c.

But now let us focus on the case of Felix who has 1025 hairs on his head. Since having 1025 hairs on one's head is not a clear instantiation of the property of baldness, condition iii will obviously fail to hold in this case: speakers of c cannot be entirely confident that 'is definitely bald' applies to Felix. Nor can they be entirely confident that it does not apply to Felix. As a result of the arising uncertainty, the statement "Felix is definitely bald" will occupy a borderline status in c. The borderline status of "Felix is bald" can then be immediately deduced by using the above inference rule. But note, again, that it would be fallacious to conclude from this that "Felix is bald" ought to be regarded as a definitely borderline statement. Uncertainty does not create a sharp demarcation between non-borderline and borderline statements. What speakers of c are supposed to be uncertain of is the correctness of the applicability of the predicate 'is bald' to Felix. This is not the same as to suppose that they are forced to think that the statement 'Felix is bald' is definitely incompatible with polar truth values. What we need is exactly the opposite of that supposition. Namely, we may plausibly hold that borderline statements must be thought to be compatible with truth and falsity. Although speakers of c are in doubt whether or not 'is bald' applies to Felix, they have no reason to exclude the possibility that 'Felix is bald' may be evaluated as true. And similarly, they cannot exclude the possibility that 'Felix is bald' may be evaluated as false. If this were not so, we would have to maintain

that they are entirely confident that no polar truth value can be assigned to "Felix is bald." In this case, however, the statement "Felix is bald" would not qualify as borderline in c.

So it appears that without leaving open the possibility of its being true or false, Fa cannot be recognized as having a borderline status. Borderline statements may therefore be considered as subject to the following modal convention:

> OPEN POSSIBILITY: 'Fa' is true or 'Fa' is false.

The presence of disjunction in OPEN POSSIBILITY indicates that the borderline status of Fa is compatible both with truth and falsity. This may be *prima facie* puzzling, since truth and falsity are defined not only as exhaustive but also as exclusive semantic values in the present framework. And surely, if "'Fa' is true" and "'Fa' is false" are equally compatible with the borderline status of Fa, then these statements must also be compatible with each other, which would involve a contradiction given the exclusivity of truth and falsity. What is the solution to this puzzle?

The proposal of the present paper is the following. The statements "'Fa' is true" and "'Fa' is false" can be presumed to be truth-functional in OPEN POSSIBILITY. This is correct, however, only if Fa itself has to be evaluated truth-functionally. I think there is reason for doubt. Remember that Fa counts as borderline in c because of its negative epistemic status, that is, because its truth value is not known to competent speakers of c. For the same reason, speakers of c are not in a position to know whether or not Fa is a truth-functional statement. In this situation, the most they are warranted in claiming to know is that Fa is compatible with a truth-functional evaluation. But because the truth value of Fa is, as a matter of fact, not known in c, the non-truth-functional evaluation may be taken to be epistemically privileged. Thus, when speakers of c apply the predicate F to a and a belongs to the borderline area of F-ness, the resulting statement will be non-truth-functional.[12]

Following this line of thought, the puzzle posed by OPEN POSSIBILITY can be dissolved. We must simply concede that Fa does not state a fact about how the world is. It does not state that a is in fact F. Rather, it states that a's being F is

---

[12] Saying that a particular statement is non-truth-functional does not imply, of course, that it is incorrect or nonsensical to assert it. On the assertability conditions of non-truth-functional statements in general, see Fredrik Stjernberg, "Restricting Factiveness," *Philosophical Studies* 146 (2009): 29-48, and Allan Hazlett, "The Myth of Factive Verbs," *Philosophy and Phenomenological Research* 3 (2010): 497-522.

an open epistemic possibility that speakers of *c* cannot eliminate.[13] In general, then, it seems more appropriate to suppose that borderline statements are governed by the following convention:

OPEN EPISTEMIC POSSIBILITY: *a* might be *F* or *a* might be not-*F*.

In contrast to OPEN POSSIBILITY, no truth-functional contradiction arises here: the epistemic possibilities of *a*'s being *F* and *a*'s being not-*F* are compatible with each other. Speakers of *c* do not possess enough epistemic information to decide whether or not having 1025 hairs on one's head counts as a clear instantiation of baldness. Thus, for all they know, it might turn out both that Felix is bald and that Felix is not bald. But they are not in a position to know which possibility is the actual one.

Conclusion

The approach delineated in the previous chapters allows us to preserve two widely held beliefs about predicate vagueness. On the one hand, it is maintained that there are specific cases of language use where the application of a vague predicate is definitely true or definitely false. On the other hand, it is agreed that in borderline cases the truth value of a vague statement remains unknown even to otherwise competent speakers. The suggested explanation for the latter fact is the non-truth-functionality of borderline statements. In using such statements, speakers do not aim at gaining or expressing pieces of factual knowledge: what they are actually aiming at is only potential knowledge. In this sense, the present approach can be regarded as a third possibility view of vagueness.

But how can a third possibility view of vague predicates do justice to the seamlessness of the sorites transitions? The most serious obstacle is removed, because no tripartite division is posited at the level of semantic values. There are only true and false application cases of vague predicates and, in addition, there are cases that resist the classical truth-functional evaluation. And because every borderline statement may be conceived as epistemically compatible with both

---

[13] There is an important analogy to the theory of epistemic modality here. According to some versions of the theory, epistemic modal statements do not serve to describe the state of affairs of the world. It is held, therefore, that no truth-functional semantics, bivalent or otherwise, is adequate for representing them. See, for example, Seth Yalcin, "Epistemic Modals," *Mind* 116 (2007): 983–1026, and Eric Swanson, "How Not to Theorize about the Language of Subjective Uncertainty," forthcoming in *Epistemic Modality*, eds. Andy Egan and Brian Weatherson (Oxford: Oxford University Press, 2011).

polar truth values, no sharp boundary can be drawn between the adjacent statements in a sorites series. Or, with other words, there is no borderline statement which can be known to be definitely incompatible with one of the two polar truth values. This is enough to rule out the existence of sharp transitions in typical sorites series.

The present theoretical approach has two further advantages. First, the non-truth-functional status of borderline statements seems to be in full accordance with our ordinary epistemic convictions. When a is recognized as a borderline case of F-ness, we are usually aware that we cannot acquire sufficient evidence or warrant to decide the question whether or not a is F. That is why we are inclined to regard the debate about such cases as unresolvable. Second, the approach also corresponds to ordinary ontic intuitions regarding the distribution of properties as they actually are in the material world. When we say that borderline cases cannot be clearly defined in terms of contrasting clear cases, we just say what we think in ordinary situations of reasoning, namely that being a borderline case of a particular property presents itself as an elusive phenomenon.[14]