DIANA RAFFMAN

VAGUENESS AND CONTEXT-RELATIVITY

In this essay I shall develop some material advanced in an earlier paper, 'Vagueness Without Paradox',[1] which proposes a solution to the sorites paradox. At critical junctures I shall summarize the relevant portions of the earlier work, but given the cumulative nature of the project in which I am engaged, there is little alternative, short of reprinting the earlier paper here in its entirety, to presupposing some familiarity with it. In this I can only beg the reader's indulgence.

   What follows is essentially a status report on work in progress. My principal goal here will be to present a new piece of my view, to "lay another brick in place"; so what I say will be long on exposition and short on argument. (Of course, a plausible dissolution of the sorites paradox is partial vindication in itself.) At many junctures the details remain to be chiseled out, and supplying those, along with the arguments needed for a full defense of my view, will require a series of further essays. I shall also not attempt to situate my approach in the larger framework of current philosophical theorizing about vague predicates, to which it frequently stands in opposition. Doing that too will be an important and illuminating task, but one for another occasion.

I

I begin by calling attention to three data – what I take to be *facts* – about the correct application of a vague predicate.

   Datum #1. The correct application of a vague predicate varies with context. As Hans Kamp observes, "[i]t is typical of a vague predicate that what objects it is true of depends on the context in which it is used".[2] For example, $100,000 can make a person rich relative to one context (say, Haitian refugees) and non-rich relative to another (say, Arab oil sheiks). A person may be tall relative to American men but

short relative to NBA players. An object may be green when viewed in outdoor light but brown when viewed in indoor light;[3] or red when viewed against one background but orange when viewed against another. Notice also that such variability of application spawns no inconsistency: we do not say that $100,000 is both rich and non-rich – that is, not both rich and non-rich *relative to the same context*. Similarly, no logical or semantic anomaly arises from the fact that *distinct* amounts of money can make a person, respectively, rich relative to one context and non-rich relative to another, even where those amounts are marginally different.[4] For example, $100,000 makes a person rich relative to Haitian refugees while $99,999 does not make a person rich relative to Arab oil sheiks, even though the two amounts differ by just one dollar. More to the point, it would make no sense to postulate a *boundary* of any sort between the two amounts, just as it would make no sense to postulate a boundary between $100,000 and itself, as it were, in the first example just offered. It would make no sense to postulate a boundary between $99,999 and $100,000 because $99,999 too makes a person rich relative to Haitian refugees, and $100,000 doesn't make a person rich relative to Arab oil sheiks either. A boundary is a division, between instances of incompatible kinds, *relative to a single context*.

Datum #2. On any run of applications of a vague predicate 'P' to the items in a sorites series, the competent speaker must reach a point at which he refuses to apply 'P'. In other words, there must be a last item to which the speaker applies 'P', and consequently a first something-other-than-P item, on any run along the series. Otherwise he is incompetent.

Datum #3. Even where the relevant context has been fixed, the location of the last item to which 'P' is applied in a sorites series will vary from speaker to speaker and run to run.[5] I may say that a given (reddish orangeish) patch is red while you say it isn't; or on one run you may say $100,000 is rich while on another run you say it isn't; and so forth. Such variation is characteristic of the *correct* application of the predicates in question. We do not say of a (competent) speaker that he was mistaken in any of these various applications; nor would you say *I* was wrong if we disagreed about the color of that patch. It is difficult to exaggerate the importance

of this fact: the correct application of a vague predicate varies from speaker to speaker and time to time.

I turn now to the sorites paradox and its implications for the correct application of a vague predicate. I shall argue that our puzzlement with the sorites results from an impoverished conception of the manner in which the items in the series can differ in kind. We tend to assume that such differences establish boundaries – differences in (incompatible) kind relative to single contexts. And we tend to assume *that* because we wrongly assume that a sorites series constitutes a single context. It does constitute what I'll call a single *external* context, but further psychological or *internal* contextual differences obtain even relative to a single series. What I say will be guided throughout by Data #1–#3.

II

To set the stage for our treatment of the paradox, consider first the verbal behavior of a competent speaker on a particular run of applications of a vague predicate to the members of a sorites series. Suppose the speaker is confronted with a series of 50 colored patches progressing from a clear case of red to a clear case of orange, so ordered that each patch is marginally different (i.e., either indistinguishable or just noticeably different) in color from the next. Suppose also that the relevant context has been fixed: for example, the entire series is to be presented at all times, in standard daylight, against a white background; the speaker is to proceed *seriatim* starting with patch #1, assessing one patch at a time, and so forth. As per Datum #2, for any given predicate, there must be a last patch to which he assigns it. (As per Datum #3, just where this last patch occurs in the series will vary from run to run.[6]) To keep things simple, suppose that on this run the speaker applies 'red' up through patch #26 and then shifts to something else – 'orange', 'reddish orangeish', 'borderline', 'don't know', 'either red or orange', 'neither red nor orange', or whatever else you like – at patch #27. For east of discussion I shall run my example using 'orange' as the contrary predicate, but nothing in my story rests on this choice.

The question then is this: How can the competent speaker do what he *must* do – viz., apply incompatible color predicates to marginally different patches?

My hypothesis is that, at the moment of judging #27, the speaker undergoes a kind of *Gestalt shift* that embraces #26 (and probably some of its predecessors) as well as #27. At the moment of shift to 'orange', the speaker is disposed to judge both #26 and #27 (plus some other patches on both sides) as being orange, thereby allowing for a change in kind while preserving the effective continuity of the series. Intuitively speaking, a string of patches shift their color simultaneously, so that #26 and #27 never differ in color *at the same time*. Like the duck-rabbit and Necker cube, these patches can "go either way": they can be "seen as red" or "seen as orange" – now one way, now the other.[7] If asked to reverse direction and retrace his steps down the series toward #1, the speaker would now judge orange some patches that he previously judged red. At some point, of course, he would shift back to 'red'; for example, he might judge #26 through #19 orange, but then undergo a Gestalt shift back to 'red' at #18. And so forth.

We can envision this verbal behavior as the output of a pair of opponent homunculi in the speaker's head – one anchored to a (mental representation of a) prototypical red patch, the other to a prototypical orange one.[8] The RED homunculus has jurisdiction over patches that come early in the series, the ORANGE one over those that come late. At least some patches in the middle, however, lie in disputed territory. Here each homunculus struggles for control, struggles to extend the scope of her jurisdiction deeper into the disputed region. A gestalt-like category *shift* occurs when one homunculus overcomes the other, pushing her backward toward her anchor or prototype. At the shift to 'orange' at #27, for example, the ORANGE homunculus overcomes the RED, pushing her back down the series toward #1. If the speaker then retraces his steps down the series, as in our example, now judging #26 through #19 orange, his RED homunculus will reassert herself at patch #18 and begin to push her rival up the series toward #50; and so forth.

If this talk of category shifts and little women in the head puts you off, think of the automatic transmission in a car. While the transmission is repeatedly shifting, accommodating the driver's changes in speed, the "phenomenology" is seamless: the driver knows that something has changed – he is now traveling faster or slower than before – but the change seems to him continuous: there are no *bound-*

aries, as it were. Analogously, when the competent speaker shifts from 'red' to 'orange', he knows something has changed yet he finds the "phenomenology" continuous: there are no boundaries in the series. (A philosopher of mind might say that the discontinuities are confined to a *subpersonal* level, thus allowing seamless continuity at the *personal* level.) The analogy goes further. Once the car has shifted to a new gear, it will continue to use that gear as long as possible, even if it slows to a speed previously handled by a lower gear. For example, if the car has shifted from second to third gear at 30 mph, it will remain in third even if it slows to 25 mph, a speed previously handled by second. (Shifting gears is hard work.) Analogously, once the competent speaker has shifted from 'red' to 'orange', if asked to retrace his steps down the series he will now call some patches 'orange' that he formerly called 'red'. The "winning" homunculus always strives to maintain her control as long as she can.

Be the mental mechanics as they may, I want to suggest that such a gestalt-like category shift constitutes a change of *context* – specifically, a change of psychological or internal context: the pre-shift internal context is distinct from the post-shift one. In homuncular terms, the relative jurisdictions of RED and ORANGE have changed. Thus patch #26 is judged red relative to one internal context, while patch #27 is judged orange relative to another. Contrary to appearances, then, a run of applications along a sorites series is not made relative to a single context; it is made relative to (at least) two contexts. In light of this new kind of contextual variation, my earlier remarks, especially my characterization of Datum #1, must be revised: what is widely acknowledged, and what we have intended by our talk of contextual variation in the application of a vague predicate, is in fact variation with what I'll call *external* context. What has gone unnoticed, and what I now mean to give center stage, is variation with *internal* context.

An example from my discussion of Datum #1 provides a helpful analogy. Consider a series of dollar amounts progressing from $1 to $100 million, each amount differing from the next by one dollar. Then suppose we ask a competent speaker to judge, of each amount, whether it makes a person rich relative to Arab oil sheiks. He begins at #1, to which he applies 'not rich', and then continues *seriatim*

along the line. Since he is competent he will be disposed to assign 'not rich' up through many millions of dollar amounts. However, suppose that when he arrives at $100,000 we change the instructions; we tell him that, from $100,000 on, he is to say of each amount whether it makes a person rich relative not to Arab oil sheiks but to Haitian refugees. Since he is competent he will say 'rich' from $100,000 on. The result will be an assignment of 'not rich' to the amounts up through $99,999 and 'rich' to the amounts thereafter – that is, an assignment that appears to contain a boundary between $99,999 and $100,000. But of course the "boundary" is illusory, appearing only to one who has lost sight of the change in context. Once that change is recognized, any air of mystery evaporates.

I want to suggest that an analogous situation obtains with respect to the red/orange series described above. It's just that in the latter case the context change is internal and largely unconscious. My claim is that whenever marginally different items are assigned incompatible predicates relative to the same external context, a Gestalt-like shift has occurred so that those predicates are assigned relative to different internal contexts. Although patch #26 is the last patch called 'red' on this run, and #27 the first patch called 'orange', no boundary is installed between them because their colors are assigned relative to different internal contexts. The upshot is that our applications of a vague predicate vary with both external and internal context. Let us define a *total context* as an ordered pair consisting of an external context and an internal context. In the case described above, 'red' and 'orange' are applied to #26 and #27, respectively, relative to different total contexts. Thus the mystery dissolves.

Of course, I have barely scratched the surface of the notion of internal context. I shall have a good deal more to say about it in the course of discussing the sorites, to which we now, finally, turn our attention.

III

Thus far I have discussed only what happens on a particular run of applications of a vague predicate to the items in a sorites series. This was a good starting place because it brings to light so much of the data that must drive the construction of an adequate account

of vagueness. But the sorites, it will be observed, is a puzzle about the application of a vague predicate independently of any particular run of judgments. For example, the paradoxical argument for 'red' and 'orange' concerns the colors the patches have *simpliciter*, as it were, independently of the ascriptions of a particular speaker on a particular occasion of use:

(1) Patch #1 is red.

(2) For any n, if patch #n is red then patch #(n+1) is red.

(3) Therefore, patch #50 is red.

What I shall claim, however, is that even considered "*simpliciter*", the correct application of a vague predicate varies with the psychological states, hence with the verbal dispositions, of competent speakers. Specifically, a vague predicate applies to an object *only relative to a total context*. As a first approximation, we can express the claim formally like this:

(V) For any object O, vague predicate 'P', and total context TC: 'P' applies to O, relative to TC, just in case a competent speaker would apply 'P' to O were he to judge O in TC.

I offered several arguments for (an ancestor of) principle V in "Vagueness Without Paradox", and I shall briefly review some of those in section IV below.[9] First, though, I want to spell out my solution to the sorites, so for now I shall simply assume the truth of V.

We are working now with 50 colored patches considered independently of any token run of applications by a particular speaker. To ensure clarity on this point, let's suppose that patch #1 is in Paris, patch #2 in Florence, patch #3 in New York, and so forth; maybe they are constituents of famous paintings – one in the Louvre, one in the Ufizzi, one in the Guggenheim, etc. The idea is that there just are these patches, somewhere in the world, whose colors form a continuous progression from red to orange; no "spatial" series need be instantiated. The question is: what colors do the patches have? If #1 is red and #50 is orange, it seems there must be a boundary somewhere in the series.[10]

According to V, the patches have their colors relative to total contexts: relative to a total context TC, each patch has the color it

would be judged to have were it to be judged in TC by a competent speaker. So if we want to specify the colors of the patches, we need to specify the relevant total context. Suppose the relevant total context TC1 consists of external context EC and internal context IC1. Then each patch has the color that a competent speaker would assign it, relative to EC, were he to judge it in IC1. EC specifies, say, standard daylight, a white background, patches of a certain size, and so forth; fill in the details as you like. IC1 is a little trickier to specify. I have not yet developed a complete account of internal contexts, but some simple examples should make the idea serviceable for present purposes. Let me emphasize that I am not wedded to the details of what follows; in the final analysis things may turn out differently from what I outline here. My present objective is simply to show that the sorites, and the correct application of vague predicates generally, are susceptible to a certain *kind* of treatment.

Specifying IC1 is tricky because the notion of an internal context turns the standard conception of a context on its head in at least two ways. First, it's not so much that we are in these internal contexts as that they are *in us*. I shall often describe us in the former way, but it must be remembered that these contexts are psychological contexts, internal states of our mind-brains. Second, on the view I propose, internal context frequently functions as a *variable*. Our normal practice is to hold context fixed and ask about the value of some variable as a function of the context in question. In respect of internal context, however, we must ask also about the context a speaker is in as a function of ... well, of something else. Let me explain.

An internal context is a psychological state of a competent speaker that grounds his dispositions to apply certain predicates in certain ways. I am not sure how finely internal contexts will need to be individuated (see notes 11 and 12 below), but we can begin by grouping them coarsely according to pairs of incompatible predicates that permit the construction of sorites series – like 'red' and 'orange', for instance. ('Red' and 'reddish orange', 'red' and 'borderline', 'orange' and 'no fact of the matter', etc., would do as well.[11]) In the terms of our homuncular metaphor, a RED/ORANGE internal context will be a state of the competent speaker's RED and ORANGE homunculi – in particular, a state of their respective jurisdictions –

that disposes him to apply 'red' and 'orange' in certain ways. As long as his dispositions remain the same, he is in the same internal context; a change in his dispositions reflects a shift to a new RED/ORANGE context. As a first approximation, then, internal contexts are identical insofar as they dispose the speaker to make the same applications of the relevant predicates, different otherwise.[12]

Exactly what dispositions are we talking about here? In particular, what dispositions are had by a speaker in internal context IC1? I shall need to approach this question in a somewhat roundabout way.

Consider a competent speaker in some initial RED/ORANGE context or other – you yourself at this very moment, say. (You are always in some RED/ORANGE internal context or other.) Let this initial internal context be IC1. Now were you to be asked the color of each patch in the series, relative to EC, you would make some assignment or other. Suppose that, were you now asked the color of patch #1, relative to EC, you would judge it red. Were you now asked the color of patch #2, relative to EC, you would judge it red. Similarly for each of patches #1–#26. (*N.B.* Here you would not be judging the patches *seriatim*. Each counterfactual situation involves the judgment of a single patch relative to EC; so the order in which your counterfactual judgements are specified is immaterial.) Of course, since you are competent, you would have to assign 'orange' to at least some of the patches – to #50, if nothing else. Keeping our example as simple as possible, let us suppose that, were you now asked the color of any of patches #27–#50, relative to EC, you would judge it orange.

Call the resulting pattern of counterfactual assignments ('red' to #1–#26 and 'orange' to #27–#50) a *complete assignment* of colors to the patches in the series, relative to EC – complete because it assigns a color to every patch. Clearly many such complete assignments exist relative to EC: different assignments generated by different speakers and by the same speaker at different times. Equally clearly, any complete assignment of color predicates to the patches will assign incompatible predicates to the members of at least one pair of adjacent (marginally different) patches. In general, any complete assignment to the members of a sorites series will assign incompatible predicates to the members of at least one pair of adjacent items.

The complete assignment just described (call it 'CA') tells us something about your verbal dispositions in your initial internal context IC1. We must proceed carefully here, however. Among other things, CA *seems* to indicate that in IC1 you are disposed to judge #1–#26 red and #27–#50 orange. But according to my view, no complete assignment to the items in a sorites series can obtain relative to any single internal context; in particular, the counterfactual judgments that determine CA could not be made entirely in IC1. At the very least, you would have to shift out of IC1 either upon considering #26 or upon considering #27. We know this because we know from our discussion in Section II that no single internal context can dispose a competent speaker to ascribe incompatible predicates to marginally different items; therefore one and the same internal context cannot dispose a speaker to judge #26 red *and* to judge #27 orange. As I have explained, wherever marginally different items satisfy incompatible predicates relative to a single external context, a Gestalt-like shift must have occurred so that those predicates apply relative to different contexts. This will be true whether the patches are judged singly, as in the counterfactual situations just described, or *seriatim* on a particular occasion, as in the scenario of section II above.[13]

At a minimum, then, in your initial context IC1 you are disposed not simply to judge #26 red and #27 orange, but to judge at least one of them in a context distinct from IC1 itself. To put the point another way, whereas all of the patches can be *considered* in IC1, not all of them can be *judged* in IC1. We are already, if unwittingly, familiar with this distinction. Recall the scenario of section II above, in which a shift to a new context was triggered by the speaker's consideration of #27. Upon considering #27, the speaker shifted out of the context he had been in after judging #26, and into a new context in which he now judged #27 orange. I shall say that the speaker *considered* #27 in one context but *judged* it, i.e., assigned it a color, in another context; and the fact that the shift to the new context occurred upon consideration of #27, rather than of (say) #26 or #28, was partly a function of the context he was in when he considered #27. Here we see clearly how internal context behaves as a variable: the context of judgment varies as a function of the context of consideration and the patch being considered, among other things.

To interpret CA correctly, then, we must recognize that it tells us what context the patches are *considered* in, not what context they are or would be *judged* in. In particular, CA tells us only what you initial context IC1 is, i.e., the context in which you initially consider the patches: it is that context in which your (counterfactual) consideration of the patches yields the complete assignment CA. (In homuncular terms, IC1 is the context in which RED has jurisdiction over #1–#26 and ORANGE over #27–#50.) In other words, CA tells us that you are initially in a context such that were you to consider any of patches #1–#26, you would judge it red, and were you to consider any of patches #27–#50, you would judge it orange. It tells us nothing, or anyway very little, about the context(s) in which those judgments would be made. Considering is not judging. Indeed, we know that either #26 or #27 could not be judged in IC1, even if that were the context in which it was considered.

Well, which of the two, #26 or #27, could not be judged in IC1? The answer is straightforward. The patch whose consideration would trigger a shift out of IC1 is the patch which is such that, having judged it, you would now generate a complete assignment different from CA. In other words, if your complete assignment is the same both before and after judging a given patch, then no contextual shift has occurred. We can put the point by saying that the patches that can be judged in IC1 are those whose context of consideration coincides with their context of judgment. In homuncular terms: the patches that can be judged in IC1 are those whose consideration would leave the relative jurisdictions of RED and ORANGE unchanged. Thus the device of the complete assignment serves as a probe for internal context change. Which of the two patches would trigger a shift out of IC1 is an empirical question. Let's suppose it is #27. This means that your complete assignment after judging #27 would diverge from CA. Perhaps your new complete assignment would assign 'red' to #1–#18 and 'orange' to #19–#50. In any even it would reflect a triumph of ORANGE over RED.

What about #28–#50? Would consideration of any of these patches trigger a shift out of IC1? Again, an empirical question. Just for example, given certain familiar "order effects",[14] consideration of a clearly orange patch like #50 might tend to "strengthen" the ORANGE homunculus, thereby enabling her to push her RED oppo-

nent backward toward #1 and establish a new internal context. Any number of scenarios are possible. But I want to keep things simple. So let's assume, perhaps implausibly, that #27 is the only patch in the series that would trigger a shift to a new internal context, and that each of #28–#50 would be judged orange in IC1.

We can now characterize IC1 more fully. IC1 is that internal state of a competent speaker in virtue of which he is disposed to judge #1–#26 red, to judge #28–#50 orange, and to judge #27 orange in some new internal context or other. More explicitly: IC1 is that state in which the speaker is disposed to judge #1–#26 red in IC1, to judge #28–#50 orange in IC1, and to judge #27 orange in some new context or other. Once in this new context (call it 'IC2'), he will enjoy a new set of dispositions; the new context will constitute a new context of consideration for the patches in the series. For instance, he might now be disposed to judge #1–#18 red and #30–#50 orange, and to judge #19–#29 orange in yet another new context.

Finally we can say what colors the patches have relative to TC1. According to principle V, the patches have the colors a competent speaker would assign to them were he to judge them in TC1. In other words, a patch has its color(s) relative to contexts in which it can be *judged*, not relative to contexts in which it can only be *considered*. Principle V should be made explicit thus:

(V)   For any object O, vague predicate 'P', and total context TC: 'P' applies to O, relative to TC, just in case a competent speaker could judge O in TC and, were he to judge it in TC, he would apply 'P' to it.

Relative to TC1, then, patches #1–#26 are red and patches #28–#50 are orange: a competent speaker could judge them in TC1 and, were he to judge them in TC1, he would apply 'red' to #1–#26 and 'orange' to #28–#50. What about #27? #27 cannot be judged by a competent speaker in IC1. A competent speaker in IC1 would judge #27 *in a different internal context*; more precisely, in IC1 he would have a disposition to judge #27 orange in another context. Hence relative to TC1, #27 has no color. Call #27 a "no-status" patch, relative to TC1. (Of course, #27 is orange relative to IC2 among others.) At first blush the idea of no-status items seems bizarre, but we shall see that it proves innocuous.[15]

Relative to any given total context, then, a sorites series contains three kinds of items: those that satisfy 'P' (e.g., 'red'), those that satisfy the relevant contrary of 'P' (e.g., 'orange'), and those that have no status (e.g., no color). Thus there will be no single total context relative to which all of the items in a sorites series have a status.[16] The latter claim appears stronger, hence less plausible, than it actually is, because it is misread as a claim about merely *external* contexts. As I remarked above, when in the past we have acknowledged the context-relativity of the application of a vague predicate, the notion of context we had in mind was that of an external context. After all, the notion of an internal context is newly introduced, so it can't be *that* sort of context we've had in mind. What we have previously thought, I submit, is that relative to any given external context, each patch has the color it would be assigned in that context by a competent speaker.[17] And of course relative to any given external context there will be some, indeed many, complete assignments of colors to the patches: at any given time a speaker would, if queried, assign some color of other to each patch.

The point is that what we have hitherto had in mind when we asked about the colors of the patches relative to a given context was always some complete assignment, i.e., some assignment of colors to all of the patches in the series relative to a given external context. And of course any such assignment will create the illusion of a boundary between adjacent patches; that is, it will appear to contain some #n such that #n is red and #(n+1) is something incompatible with red, *relative to a single context*. But that "single" context is only an external context, and so the seeming boundary vanishes in light of further, internal, contextual variation. In considering only external context, we have omitted a crucial element of the story.

The mistake is a natural one since we competent speakers never know what internal context anyone is in. All we can know, indeed all that ever interests us, is what the operative external context is. All we can do, when we want to know the colors of objects, is select the appropriate external context, put ourselves in it, and see what we say. Then we can know that, relative to that external context, the objects in question have those colors in some internal context(s) or other. Often we can know that the internal contexts involved must be different, as when we judge patch #26 red and patch #27 orange,

or I say that a 17-week-old fetus is a person and you say it isn't; but we can't know *which* contexts they are. We never need to know that, though, so things work out fine.

I said that the idea of no-status items is innocuous. It seems strange only because we fail to appreciate how little is being claimed. In particular, there is no time at which any patch or other visible object lacks a color: for every visible object, there is at any time some total context(s) relative to which it is colored (see again note 17). Simply put: every visible object is colored at all times. Thus we competent speakers never *encounter* a no-status item. "No-status" is not an assignment ever made by a competent speaker, because the very act of considering a patch or other object with respect to color always occurs in some internal context or other. Internal contexts travel with us in our heads, as it were, and so we are always, automatically, *in* the requisite context. As with the automatic transmissions in our cars, we are for all intents and purposes unaware of those contexts and the shifts between them. Indeed, *we* don't make the shifts; some of our subpersonal parts do. Any time we go to assess the color of an object relative to a given external context – and that much *is* something *we* do – an internal context is automatically and tacitly established for us.

In the end, the sorites dissolves because its major premise is false. Relative to any total context, there will be some #n such that #n is P and #(n+1) has no status (e.g., #n is red and #(n+1) has no color). Alternatively, we can think of the situation this way: for any complete assignment made relative to a given external context, there will be some #n such that #n is P, relative to one internal context, and #(n+1) is something incompatible with P, relative to another internal context.[18]

Plainly there remain a great many questions to be answered, details to be fleshed out, worries to be quelled. I have here tried simply to take a further step toward an adequate account of vague predicates, and in particular to explicate the role of psychological context-relativity in their application. I close now with a brief consideration of some remaining questions.

IV

Aren't there now boundaries between the red and no-status patches, on the one hand, and between the no-status and orange patches, on the other, relative to any given total context? Doesn't this simply reintroduce the original problem?

Here I think we can say what we like. On the one hand we can simply deny that the divisions in question *are* boundaries, on the ground that they don't constitute differences in (incompatible) kind between marginally different items relative to a single (total) context. Again, 'no-status' does not name a property on a par with red and orange; rather, it signifies the absence of any color. On the other hand, we can say that there are such boundaries but that they are not boundaries of a sort that should worry anyone. In the first place, as I observed above, we competent speakers never *encounter* any no-status items, hence never encounter the boundaries in question. Second, the divisions in question are not semantically legislative; they carry no normative force, underwrite no distinction between correct and incorrect usage. A speaker for whom the last red patch is, say, #26 or #28 instead of #27 is not therein mistaken. Rather, these divisions reflect Gestalt-like contextual shifts of a purely mechanical (i.e., psychological) nature. Their location is in this sense arbitrary. More to the point, there is no reason to shift, hence no justification for shifting, at any particular #n in the series as opposed to #(n−1) or #(n+1). If there were a reason, then either the predicate in question would not be vague or the differences between adjacent items in the series would not be marginal in the sense required to generate a paradox. Third, the present appeal to intuition is in any case illegitimate. The notions of internal context and no-status items are newly introduced, so it is hard to see how we could have any legislative *intuition* – that is, any intuition that should constrain our theorizing – to the effect that the boundaries in question are problematic.

Finally, as promised, I shall review briefly some of the arguments presented in "VP" for the principle *V* (see again note 9 below). Specifically, these are arguments for the claim that the application of a vague predicate varies with psychological or internal context.

First, if the correct application of a predicate does *not* vary in something like the ways I describe here, then it is hard to see how the

predicate can be vague. Among other things, if the correct extension of the predicate were invariant with internal context, then once the relevant external context had been fixed, the shift to a new predicate (e.g., 'orange') would occur at the same item on every run along a sorites series; in other words, there would be sharp and fixed boundary in the series and the predicates in question – or anyway the best use we could make of them – would not be vague. Similarly, if a predicate's extension were invariant with internal context, there would be a real boundary, not merely an apparent one, between adjacent items in any complete assignment. Or so I have been at pains to show.

Second, since our actual applications of vague predicates do appear to vary with psychological context, it follows that if the *correct* application of these predicates does not vary in this manner, then we are linguistically incompetent in their use: we do not – presumably cannot – know the correct extensions of the predicates at issue. But we are *not* incompetent in their use. Incompetence with such a vast portion of our natural language – indeed our entire natural language, some would contend – is not a credible hypothesis. Granted, we cannot know, in advance of our use of a vague predicate on a particular occasion, what judgments we will make, where our category shift(s) will occur, or what our new context(s) will be. But that is a far cry from claiming that we can never know the correct extension of a vague predicate.

In addition to the foregoing arguments, and perhaps more compelling than any of them, is the independent plausibility of the idea that the correct application of a vague predicate varies with the judgments of competent speakers. Compare 'precisely $100,000' and 'roughly $100,000', 'contains precisely 1,112 grains of sand' and 'is a heap of sand', 'is precisely 17 weeks old' and 'is a person'. Surely the *raison d'être* of these vague expressions is to enable us, imprecise creatures, to negotiate the world around us. As Crispin Wright observes, "[w]e should have no use for a precisely demarcated analogue [of a vague predicate] in contexts in which the word is typically used. It would, for example, be ridiculous to force the question of obedience to the command 'pour out a heap of sand here', to turn on a count of the grains" ("LM", p. 231). How could vague

words play their intended proprietary role, if not by a dependence upon the judgments of competent speakers?[19]

NOTES

[1] *Philosophical Review*, vol. 103, no. 1, 1994: 41–74 (hereafter, "VP"). See also my "Comments on Terence Horgan's *Transvaluationism*", *The Southern Journal of Philosophy*, vol. XXXIII, Supplement, *Proceedings of the 1994 Spindel Conference*: 127–132. The reader will here notice several revisions to the earlier view.

[2] "The Paradox of the Heap", in *Aspects of Philosophical Logic*, ed. U. Monnich (Dordrecht: D. Reidel): 225–228.

[3] Even where both kinds of light count as normal; see C. L. Hardin, *Color for Philosophers* (Indianapolis: Hackett Publishers, 1988), p. 162.

[4] The term 'marginal difference' comes from Crispin wright; see for example "Language Mastery and the Sorites Paradoz" (hereafter, "LM"), *Truth and Meaning*, ed. Gareth Evans and John McDowell (Oxford: Oxford University Press, 1976): 223–247.

[5] See "VP", pp. 51–52.

[6] A great deal depends on the order in which he judges the patches; see "VP", p. 52.

[7] Again, see "VP" for the full story. The idea of Gestalt-like shifts may be more immediately plausible for observational words like 'red' and 'bald' than for non-observational ones like 'rich' or 'person'; but if I'm right, the analogy works across the board. It's just that whereas a shift from 'red' to 'orange' is a shift in *perceptual* aspect, a shift from 'rich' to 'non-rich', or from 'person' to 'non-person', is a shift in *conceptual* aspect.

[8] The homuncular scenario just sketched differs from the one found in "VP", pp. 47–49.

[9] See "VP", pp. 65–66. The principle defended in "VP", so-called *Biconditional B*, differs slightly from *V*; among other things, I had not then introduced the idea of a total context. Similar arguments should work in both cases, however.

[10] This question was posed to me by Timothy Williamson.

[11] It may well be that a RED/ORANGE internal context will need to incorporate the relevant specifications for all of these predicates in the 'red'-to-'orange' range. In that case the complete story would be only more complicated than what I outline here, not essentially different.

[12] It may be that internal contexts are identical only if the dispositions they ground are identical across all external contexts. The niceties of individuation remain to be worked out.

[13] If anything, the justification for positing a shift is greater in the former case. Absent a shift there, the speaker would be disposed to apply incompatible predicates to marginally different (possibly indistinguishable) items while in the same *brain* state.

[14] "VP", pp. 51–52.

[15] Though I cannot pursue this issue here, let me ward off a potential confusion. It does not follow from #27's lack of a color relative to TC1 that the sentence

'#27 is red' has no truth-value or an "indeterminate" truth-value relative to TC1. Rather, that sentence is false: an object that has no color is not red. Impressions to the contrary result from a misconception of the nature of borderline cases and of the incompatible predicates involved in any sorites paradox. See my "Borderline Cases and Excluded Middle: A Common Misconception" (ms, in progress; hereafter, "BC") for detailed discussion.

[16] Do not confuse no-status items with borderline cases. I cannot give this important distinction its due here, but let me say the following. The notion of a borderline case is ambiguous as between what I call its 'object-linguistic' and 'metalinguistic' senses. The object-linguistic sense is the one that invites confusion with the idea of a no-status item, and on this score my view is as follows. A borderline case (in the object-linguistic sense) has a color that joins with red and orange *inter alia* to form a family of incompatibles. In other words, borderline cases do have a "status": they have a color, namely borderline red or borderline orange (*a.k.a.* red-orange? reddish-orangeish?), that is neither red nor orange but somewhere between the two; and relative to any given total context TC, a patch is a borderline case if and only if a competent speaker would judge it a borderline case were he to judge it in TC. No-status patches, on the other hand, have no color relative to TC. See my "BC" for more.

[17] We can talk this way provided we are careful. I am not saying that objects have *no* colors relative to any external context, but rather that they have their colors relative to *more* than external context. It may help to keep in mind that, relative to any given external context, all of the patches will have some color and most patches will have more than one; whereas relative to any given total context, only some patches will have colors and no patch will have more than one. (See "VP", pp. 62–63.)

[18] Of course, any adequate treatment of vague predicates must explain the powerful intuition that the major premise of the paradoxical argument is true; see "VP", pp. 44–58, for detailed discussion. In "VP" I argued that the major premise was ambiguous as between (false) "singular" and (true) "pairwise" readings. It now seems to me that the charge of ambiguity was ill-conceived. The major premise is simply a false claim about the items in the series considered singly. We *think* it is true because we mistake it for a different claim (what I earlier called 'IP*') about the items considered pairwise. I thank Richard Cartwright for pressing me on this point.

[19] Robert Batterman, Mark Richard, Stewart Shapiro, Barbara Scholz, and Don Hubin are responsible for any errors in this paper.

Department of Philosophy
Ohio State University
Columbus, OH 43210-1365
USA