

‘I don’t know what you mean by “glory”’, Alice said.

Humpty Dumpty smiled contemptuously. ‘Of course you don’t—till I tell you. I meant “there’s a nice knock-down argument for you!”’

‘But “glory” doesn’t mean “a nice knock-down argument”’, Alice objected.

‘When *I* use a word’, Humpty Dumpty said in rather a scornful tone, ‘it means just what I choose it to mean—neither more nor less’.

‘The question is’, said Alice, ‘whether you *can* make words mean so many different things’.

‘The question is’, said Humpty Dumpty, ‘which is to be master—that’s all’.

Lewis Carroll, *Through the looking glass*

# Vagueness and Conversation

Stewart Shapiro

The Ohio State University

The University of St. Andrews

**1 Determinacy.** The present account of how vague terms arise in natural language, and how their meaning is determined, begins like that of Vann McGee and Brian McLaughlin (1994) (see also McGee 1991). It is a truism that the competent users of a language somehow determine the meaning of its words and phrases, by what they say, think, and do. Lewis Carroll’s Humpty Dumpty says that it is a question of ‘which is to be master’. When it comes to meaning, the proverbial *we*—the community of competent, sincere, honest, well-functioning speakers—are master. To be sure, I have no special

insights to bring as to how words acquire their meaning. There is also a nice question concerning who, exactly, are the competent speakers of a natural language like English. Or, perhaps better, what is it to be a (competent) speaker of English? Who are ‘we’?

Let  $F$  be a monadic (possibly complex) predicate in a natural language. McGee and McLaughlin (1994, §2) introduce a technical term ‘definitely’ as follows:

... to say that an object  $a$  is definitely an  $F$  means that the thoughts and practices of speakers of the language determine conditions of application for ...  $F$ , and the facts about  $a$  determine that these conditions are met.

In a note, they add that the word ‘determinately’ would be better, since it captures the underlying idea. I will use the latter term here. In general, when ‘a sentence is [determinately] true, our thoughts and practices in using the language have established truth conditions for the sentence, and the [presumably] non-linguistic facts have determined that these conditions are met’. McGee and McLaughlin insist that the ‘determinately’ operator is not, or may not be, compositional. It may be that a sentence of the form ‘ $\Phi \vee \Psi$ ’ is determinately true even if it is not the case that  $\Phi$  is determinately true, nor is it the case that  $\Psi$  is determinately true. More notably, it may be that a sentence  $\Phi$  fails to be determinately true without  $\neg\Phi$  being determinately true. With McGee and McLaughlin, I believe that this is the source of vagueness. An object  $a$  is a *borderline* case of a predicate  $F$  if  $Fa$  is ‘unsettled’, i.e., if  $a$  is not determinately an  $F$ , nor is  $a$  determinately non- $F$ .

Since we are still at the level of truisms and definitions, nothing should be all that controversial so far. McGee and McLaughlin are free to define technical terms as it pleases them, and I am free to follow suit. One point of contention is whether or not ‘determinately’, so defined, is well-defined. I will have a bit more to say on this by way of further articulation (e.g., §6 below), but will just assume for now that we have hold of at least a vague notion here, which can give us a start on theorizing. Another point

of contention is the thesis determinateness is not compositional. I do not know how to defend this, and so will take it as a premise toward what follows.

A more crucial point of contention is whether there are any borderline cases at all—as the term ‘borderline’ is defined here (in terms of the technical notion of determinateness). Famously, Timothy Williamson’s (1994) epistemicism maintains that the thoughts and practices of language users fix precise extensions, and complementary anti-extensions, for every predicate of the language. Consider, for example, the sentence ‘Josh arrived around noon’. According to Williamson, the established use of language determines a real number  $n$  such that if Josh arrived exactly  $n$  seconds after noon, then the sentence is true, but if he arrived  $n+.0000001$  seconds after noon, the sentence is false. Williamson would thus have little use for the above definition of ‘borderline’, since on his view, there simply are no ‘unsettled’ cases. Williamson defines the word ‘borderline’ in terms of what can be known, rather than in terms of truth conditions, as above. Here, of course, I insist on the above definition, noting that the property might be empty (but hoping that it isn’t).

**2 Tolerance and open-texture.** Crispin Wright (1976, §2) defines a predicate  $F$  to be *tolerant* with respect to a concept  $\phi$ , ‘if there is . . . some positive degree of change in respect of  $\phi$  insufficient ever to affect the justice with which  $F$  applies to a particular case’. Vague predicates, or at least some vague predicates, are said to be tolerant. Wright’s language suggests that the issue concerns proper *judgement*, which I take to be an important insight.

Consider the predicate ‘bald’, which is arguably tolerant with respect to number and arrangement of head hair. Suppose that two men  $h, h'$  differ only marginally in the amount and arrangements of their head hair. Wright’s principle seems to say that if someone competently judges  $h$  to be bald, then she must judge  $h'$  to be bald too—either both are bald or neither are. Similarly, the predicate ‘red’ is tolerant with respect to indistinguishable, and perhaps small, differences in color. If two colored patches  $p, p'$  are

visually indistinguishable, and if someone judges  $p$  to be red, she must judge  $p'$  to be red as well. As plausible as this interpretation of tolerance may seem, especially in the case of the colored patches, it leads to absurdity in the now familiar manner. In general, a sorites series arises when we have a (prima facie) tolerant predicate  $P$ , and a series of objects running from a clear (or definite) instance of  $P$  to a clear non-instance of  $P$ , with each differing marginally from its neighbors.

Perhaps we can agree on the following weaker *principle of tolerance*:

Suppose a predicate  $P$  is tolerant, and that two objects  $a, a'$  in the field of  $P$  differ only marginally in the relevant respect (on which  $P$  is tolerant). Then if one competently judges  $a$  to have  $P$ , then she cannot judge  $a'$  to not have  $P$ .

This, I submit, is the key to avoiding contradiction. Suppose that  $h$  and  $h'$  differ only marginally in the amount and arrangements of their head-hair. It is compatible with the principle of tolerance that someone judge  $h$  to be bald and leave the bald-state of  $h'$  unjudged (one way or the other). It does violate tolerance, in the present sense, if a subject judges  $h$  to be bald and *decides* to leave  $h'$  unjudged. The observation here is that the subject does not violate tolerance if he has not considered the state of  $h'$ . The principle of tolerance demands that if this person were then asked, or forced, to judge  $h'$ , and *does not change his mind* about  $h$ , then he must judge  $h'$  to be bald. Suppose that our subject judges  $h$  to be bald and we inquire about  $h'$ . Then he can satisfy the principle of tolerance by either judging  $h'$  to be bald, or by judging  $h'$  to be not bald and taking back his previous judgement that  $h$  is bald. I suggest that in the cases of interest, the meaning of the word 'bald' and the semantic and non-semantic facts allow this option.

Let  $P$  be a tolerant predicate for which there is a sorites series: a list  $a_1, \dots, a_n$ , where it is determinate that  $Pa_1$ , it is determinate that  $\neg Pa_n$ , and for each  $i < n$ , the difference between  $a_i$  and  $a_{i+1}$  is insufficient to affect the justice with which  $P$  applies. As noted above, I take it as a premise (without much in the way of argument) that somewhere in the series there is a borderline case of  $P$ . There is a  $j$

such that  $Pa_j$  is neither determinately true nor determinately false. Since  $P$  is a predicate of a natural language like English, and  $a_j$  is in its range of applicability, our ‘thoughts and practices in using the language have established truth conditions for’ the sentence  $Pa_j$ . My premise is that despite this, the (presumably) non-linguistic facts have not determined that these conditions are met. Moreover, language users have established truth conditions for  $\neg Pa_j$ , and those conditions have not been met either. As far as the language has evolved to date,  $Pa_j$  is still open. As McGee and McLaughlin put it, the sentence is ‘unsettled’.

To be sure, this is not to say that a given competent speaker, when asked about  $Pa_j$  (under normal conditions, whatever those are) realizes that the sentence is indeterminate, or has the phenomenological feel that there is nothing to say or that she can go either way. Some competent speakers may feel that way, but other competent speakers may be inclined to assert  $Pa_j$ , while still others might be inclined to assert  $\neg Pa_j$ —all under the same circumstances. The claim now is only that if  $Pa_j$  is borderline, then the meanings of the words, the non-linguistic facts, etc. do not determine a verdict, even if some individual language users do themselves competently assert a verdict in some circumstances.

The presence of ‘unsettled’ cases is precisely the sort of situation that suggests a supervaluational semantics, much like the one provided by McGee and McLaughlin (and Fine 1975). The idea is to introduce a notion of an admissible *sharpening*, or precisification, of the language. In the case at hand, a sharpening of the language would contain a sharp (i.e., non-vague) predicate  $P^*$  such that for all  $x$  in the range of  $P$ , if  $Px$  is determinately true, then  $P^*x$ , and if  $\neg Px$  is determinately true, then  $\neg P^*x$ . A sentence  $\Phi$  is *super-true* if it comes out true under all admissible sharpenings of the language. So super-truth coincides with determinate truth, as that notion is defined above.

I agree that the overall supervaluational framework is natural and helpful here—eventually—but there are features of the use (and thus the meaning) of vague terms that need to be incorporated into it.

Since we encounter, and effectively and competently decide, borderline cases regularly, there is more to truth than super-truth, and there is more to validity than the necessary preservation of super-truth.

Suppose, again, that  $a$  is a borderline case of  $P$ . I take it as another premise that, in at least some situations, a speaker is free to assert  $Pa$  and free to assert  $\neg Pa$ , without offending against the meanings of the terms, or against any other rule of language use. Unsettled entails open. The rules of language use, as they are fixed by what we say and do, allow someone to go either way in the borderline region. Let us call this the *open-texture* thesis.

Wright (1987, 244) seems to endorse the open-texture thesis (at least parenthetically):

‘Borderline cases are . . . cases about which competent speakers are allowed to differ.’ And Sainsbury (1990, §9): ‘Given the nature of boundarylessness, semantics give freedom. There is some number of minutes such that the nature of the concept of a person, together with the nature of the world, makes it neither mandatory nor impermissible to apply the concept to a foetus of that age in minutes.’

Open-texture is a more or less empirical claim about the proper use of vague terms in language. However, I have very little in the way of empirical evidence to offer for it from this armchair, beyond a few intuitions. That’s why I call it a ‘premise’. I realize that with the assertion of open-texture, I am leaving the comfort of consensus and truism (if it has not been left already). Some supervaluationists claim that if  $a$  is a borderline case of  $P$ , then  $a$  lies *outside* the extension of  $P$ , and so it is (determinately) incorrect to assert  $Pa$ . It is a platitude that one should assert only truths (of which, see §6 below). And, again, the supervaluationist idea is that truth is super-truth. Since, by hypothesis,  $Pa$  is not super-true, it is not correct to assert it—or so argue these supervaluationists. Similarly, those inclined toward a fuzzy approach claim that if  $a$  is a borderline case of  $P$ , then  $Pa$  is less than completely true. Since, strictly speaking, one should assert only full truths, these theorists hold that in strict circumstances, it is not correct to assert  $Pa$ , in agreement with our supervaluationists. The fuzzy logicians might add that even in less strict situations, one should assert only *near* truths, and some borderline cases are not near truths. In

contrast, I would take the premise of open-texture to apply generally, to any borderline case, no matter how ‘far’ it is from a determinate one. I will not pause to offer arguments against these opposing views, or the intuitions that lie behind them, but will rest content to let the entire framework serve as a rival account.

Let’s be clear as to what the open-texture premise is, and what it is not. Suppose that  $a$  is a borderline case of  $P$ . As above, the open-texture thesis does not entail that a given competent speaker will always be conscious of the fact that he can go either way. The open-texture thesis is that in some circumstances, he can, in fact, go either way without offending against the meaning of the terms, the non-linguistic facts, and the like.

Second, even if  $a$  is a borderline case of  $P$ , it is not true that the rules for language use allow a speaker to assert  $Pa$  in any situation whatsoever. For example, one is not free to assert  $Pa$  if one has just asserted (and does not retract)  $\neg Pa$ . This would offend against logic (dialetheism notwithstanding). Similarly, one is not normally free to assert  $Pa$  if one has just asserted (and does not retract)  $\neg Pa'$ , where  $a'$  is only marginally different from  $a$ . That would offend against tolerance.<sup>1</sup>

In short, if a sentence  $\Phi$  is determinate, then one can correctly assert it at any time and can never correctly deny it. If  $\neg\Phi$  is determinate, then one can correctly deny  $\Phi$  at any time, and can never assert it. If  $\Phi$  is unsettled, then it depends on the conversational situation one is in. We need a mechanism to track the features of a situation that allow and disallow certain moves. For this, I turn to Lewis (1979).

---

<sup>1</sup>This is not to say that tolerance can never be violated. To follow an example from Sainsbury (1990), consider an art shop that has red paints on one shelf, marked ‘red’, and orange ones on another, labeled ‘orange’. The proprietor would not be judged incompetent concerning English if one cannot really tell the last red jar from the first orange one.

**3 Conversational score.** It is often noted that the truth values of instances of vague predicates are often relative to a comparison class. A short NBA player would not be a short business executive. A person can be wealthy in some nation, or neighborhood, while someone else with exactly the same resources is not wealthy in other situations. Similarly, the truth value of an instance of a vague predicate can vary with an instance to which it is compared. An income may be paltry with respect to Bill Gates, but not so with respect to George W. Bush. Or, to follow Graff (2000, §3), a man can be bald with respect to Yul Brynner and not so with respect to Mikhail Gorbachev (even though the comparison class is the same in both cases—all men).

The present program introduces another relativity, to a conversational context. This includes not only the environment of the conversation, what it is about, and the implicit or explicit comparison class or paradigms, but also what has already been said in the course of the conversation.

For this purpose, I borrow the notion of a *conversational score*, or conversational record, invoked by some linguists interested in pragmatics. The record, or score, is a local version of common knowledge. David Lewis's influential 'Scorekeeping in a language game' (1979) develops the idea in sufficient detail for present purposes. During a conversation, the score contains the assumptions, presuppositions, and other items implicitly or explicitly agreed to. For example, the conversational score contains the range of quantifiers like 'everyone' and the denotata of proper names like 'Barbara' and 'Joe'. It also settles the relevant comparison class and/or paradigm cases for predicates like 'tall' (NBA players, executives, etc.) and 'bald' (Brynner, Gorbachev). The score also contains propositions that have been (implicitly or explicitly) agreed to, and are not up for dispute or discussion, at least for the moment. Among these are the 'presuppositions' to the conversation.

The conversational record is a sort of running database. It is continually updated, in that items are put on the record and, notably here, removed from it in the course of a conversation. Items get removed when the topic changes, when what is agreed to or some presupposition comes into question, or



when some of the participants change their minds about items that are on the record. Since the participants in a conversation may be mistaken about what has been agreed to and what has not, they may be mistaken about what is on the conversational score. Normally, this sort of thing gets cleared up in due course, if it turns out to be important.

Lewis (1997, 345) delimits some features of the conversational database, most of which are relevant here:

(1) . . . the components of a conversational score at a given stage are abstract entities. They may not be numbers, but they are other set-theoretic constructs: sets of presupposed propositions, boundaries between permissible and impermissible courses of action, or the like.

(2) What play is correct depends on the score. Sentences depend for their truth value, or for their acceptability in other respects, on the components of the conversational score at the stage of the conversation when they are uttered . . . [T]he constituents of an uttered sentence—subsentences, names, predicates, etc.—may depend on the score for their intension or extension.

(3) Score evolves in a more-or-less rule-governed way. There are rules that specify the kinematics of score:

If at time  $t$  the conversational score is  $s$ , and if between time  $t$  and time  $t'$ , the course of the conversation is  $c$ , then at time  $t'$  . . . the score is some member of the class  $S$  of possible scores, where  $S$  is determined in some way by  $s$  and  $c$ .

(4) The conversationalists may conform to directives, or may simply desire, that they strive to steer components of the conversational score in certain directions . . .

(5) To the extent that conversational score is determined, given the history of the conversation and the rules that specify its kinematics, these rules can be regarded as constitutive rules akin to definitions.

Lewis points out that unlike most games, conversations tend to be cooperative. Indeed, that is often the point of having a conversation in the first place. '[R]ules of accommodation . . . figure prominently among the rules governing the kinematics of conversational score' (1979, 347). The idea is that the conversational record tends to evolve in such a way that, other things equal, whatever is said will be construed as to count as correct. The record will be updated to make this so, if possible. Suppose, for example, that someone utters a sentence like 'Harry no longer smokes unfiltered cigarettes', which has the presupposition that Harry used to smoke unfiltered cigarettes. Unless someone objects, the presupposition goes on the record. To be sure, cooperation is 'not inevitable, but only a tendency', as Lewis puts it. And, of course, presuppositions can be retracted later.

Lewis goes on to illuminate six diverse features of the semantics and pragmatics of natural language by using this notion of conversational score. The list includes definite descriptions, performatives, and a painfully brief treatment of our present topic, vagueness (1979, 351-354). That discussion begins by noting that with a vague predicate like 'bald', 'nothing in our use of language' fixes a sharp border between the bald and the non-bald. This much echoes the above McGee-McLaughlin theme, and like them, Lewis gestures toward the framework of supervaluation. Suppose that Fred is a borderline case of 'bald'. Then whether the sentence 'Fred is bald' is true depends on 'where you draw the line'. Relative to 'some perfectly reasonable ways of drawing a precise boundary between bald and non-bald, the sentence is true. Relative to other delineations, no less reasonable, it is false'. Philosophers and linguists cannot settle on a single, precise border, but 'must consider the entire range of reasonable delineations'. Lewis says that if 'a sentence is true over the entire range, true no matter how we draw the line, surely we are entitled to treat it simply as true'. In other words, *simple truth* is super-truth. But this is truth only in a strict sense. Lewis goes on to recognize a looser use of language:

But also we treat a sentence more or less as if it is simply [i.e., super-]true, if it is true over a large enough part of the range of delineations of vagueness. (For short, if it is *true enough*.) If a

sentence is true enough (according to our beliefs) we are willing to assert it, assent to it without qualification, . . .

He notes that we usually do not get into any trouble with this loose use of language, although we can and sometimes we do—witness sorites paradoxes.

Lewis (1979) uses the notion of conversational score to resolve an issue that concerns his notion of ‘true enough’:

When is a sentence true enough? Which are the ‘large enough’ parts of the range of delineations of vagueness? This is itself a vague matter. More important for our present purposes, it is something that depends on context. What is true enough on one occasion is not true enough on another. The standards of precision in force are different from one conversation to another, and may change in the course of a single conversation. Austin’s ‘France is hexagonal’ is a good example of a sentence that is true enough for many contexts, but not true enough for many others. Under low standards of precision it is acceptable. Raise the standards and it loses its acceptability.

To adapt Lewis’s example, suppose that someone says ‘France is hexagonal’, and gets away with it in a conversation. The rules of accommodation suggest that appropriately loose standards thereby go on the conversational record, since such standards are needed to make the assertion true (i.e., true enough). Thereafter, ‘Italy is boot shaped’ is also true enough, since that conforms to the same low standard (assuming that Italy is at least as boot shaped as France is hexagonal). But if someone later denies that Italy is boot shaped, insisting on the differences, then the standards are thereby raised, and ‘France is hexagonal’ is no longer true enough. So it comes off the record.

Lewis then provides a clean, effective reply to Peter Unger’s (1975, 65-68) argument that nothing (or hardly anything) is flat. Unger claims, and Lewis agrees, that nothing can be flatter than something that is flat. So suppose someone says that Kansas is flat. This cannot be correct, since a level sidewalk

is flatter than Kansas. But this sidewalk is not flat either, since a sheet of metal is flatter than that. And a sheet of metal is not flat, since . . . So we can only say that Euclidean planes are flat. No physical objects qualify.

Clearly, Unger's conclusion offends against language use. We use the word 'flat' in all sorts of contexts, and presumably we know what we are talking about. We also manage to communicate using the word. To echo Humpty Dumpty, who is master?

To be sure, someone who says that Kansas is flat does not mean that it is *perfectly* flat. To use Lewis's terminology, the speaker says that it is flat enough. When the remark about Kansas is made, the rules of accommodation require a relatively low standard, since Kansas does have small hills here and there (not to mention the curvature of the earth). Now suppose that later in the conversation, someone says that his sidewalk is flatter than Kansas, or complains that his sidewalk is not flat. According to Lewis, the truth of this assertion requires a higher standard for 'flat' (or 'flat enough') than had been established with the statement about Kansas. The rules of accommodation thus put a higher standard on the record. On this new standard, Kansas is not flat (enough). So once the standards are raised, 'Kansas is flat' is removed from the record.

**4 Conversation and sorites.** Unfortunately, this is the extent of the treatment of vagueness in Lewis (1979), and I am not aware of any use of conversational score with respect to vagueness in his other writings (including his extensive writing on vagueness). In particular, Lewis (1979) does not invoke conversational score to deal with sorites. Indeed, the only (possible) reference to the paradox is a remark that loose thinking sometimes leads to trouble. Given the framework he proposes, Lewis accepts the standard supervaluational resolution.

Recall our premise of open-texture: borderline cases of vague predicates can go either way in some conversational contexts. So once a borderline case is 'resolved' in the course of a conversation,

that information goes on the conversational record, and will remain on the record unless it is implicitly or explicitly retracted.

To develop a fully general account of reasoning with a conversational score, we would need a paraconsistent logic, since it is possible for conflicting propositions to be put on a record, without the group retracting any of them. The participants may not realize they have contradicted themselves. Even if they do realize the contradiction, they may not want to retract the statements, since they remain attracted to each of them individually, and do not know which one(s) to give up. Such is paradox. The presence of a contradiction on the conversational record (whether detected or not) does not commit the group to every proposition whatsoever. Nevertheless, I take it that consistency is a regulative ideal for a conversational record, in the sense that the participants try to achieve consistency, retracting questionable items from the score as needed (Priest 1987 notwithstanding). For present purposes, however, we need not invoke paraconsistency, since our focus is on admittedly artificial cases in which consistency is easily enforced.

Let us begin with a conversational version of a Horgan-style ‘forced march’ sorites situation (see Horgan 1994). Suppose we have a series of 2000 men lined up. The first is a mature Yul Brynner, who is clearly bald—he has no hair at all (or so we will assume). The last man is Jerry Garcia, in his prime, who we will take as a paradigm case of a man who is not bald. The hair of each man in the series (who has hair) is arranged in roughly the same way as his predecessor. After the first, each man differs from the one before by having only slightly more hair, perhaps imperceptibly more.

Now suppose that the participants in a conversation start asking themselves about the baldness-state of each man in the series, starting with Yul Brynner, and they insist on a communal verdict in each case. As each question in the form ‘Is man  $n$  bald?’ is put, they are to provide a ‘yes or no’ answer.

This forced-bivalence is only a convenience. The situation would be essentially the same if we gave them the option to answer ‘borderline bald’, ‘unsettled’, ‘no fact of the matter’, etc. It would also

be the same if we gave them the option to agree on silence. To paraphrase Raffman (1994, 41, n. 1), it isn't merely that there is tolerance between (determinate) cases of baldness and (determinate) cases of non-baldness. There is tolerance between baldness 'and *any* other category—even a "borderline" category'. Or Wright (1976, §1): 'no sharp distinction may be drawn between cases where it is definitely correct to apply [a vague] predicate and cases of *any* other sort'. What matters for present purposes is that the participants in our conversation must answer (or refuse to answer) by consensus—whatever the allowed answers may be. If, at any point, they want to stop answering, they must *agree* to shut up. We do not allow them to stop the march by simply failing to agree on a verdict.

So let us allow the simplifying assumption that only two answers are allowed: 'bald' and 'not bald'. Being competent speakers of English, the conversationalists all agree that Yul Brynner is bald, that the second man is bald, etc. Eventually, they will move into the borderline area, and encounter cases whose baldness state is 'unsettled', as in §1 above. Again, the thoughts and practices in using the language have established truth conditions for statements about baldness, and truth conditions for non-baldness. In the borderline region, the non-linguistic facts have not determined that either of these truth conditions are met. Nevertheless, the conversationalists in this exercise will probably continue to call the men bald as they move through the borderline area—for a while. If they call man  $n$  bald, they will probably call man  $n+1$  bald as well, since by hypothesis, they can barely tell the two heads apart in the relevant respect.

This is all right, so far as their language competence goes. Given our open-texture premise, borderline cases can go either way (without offending against meaning and the facts), and the participants in this conversation are just going one way rather than the other as they enter the borderline area from this direction. Again, it does not matter whether they realize that they can go either way in any of these borderline cases. For some of the conversationalists and for some of the cases, it certainly won't

feel like they can go either way. Following their instructions, they just call 'em as they see 'em. And as they start to move into the borderline cases, they continue to see 'em to be bald—for a bit.

This puts propositions like 'man 923 is bald' and 'man 924 is bald' on the conversational score (assuming that those are borderline cases). Now recall the principle of tolerance:

Suppose a predicate  $P$  is tolerant, and that two objects  $a, a'$  in the field of  $P$  differ only marginally in the relevant respect (on which  $P$  is tolerant). Then if one competently judges  $a$  to have  $P$ , then she cannot judge  $a'$  to not have  $P$ .

Let us assume that tolerance is in force in the present conversation. Arguably, it is part of the meaning of 'bald' that it is tolerant, and so there is usually no option to violate tolerance while being true to the meaning of the words of the language (but see note 1 above). In particular, if we were to present any *adjacent pair* from the series to the conversationalists at the same time, they would agree that either both of the men are bald or that neither are (or that both are borderline, etc., if we suspend the forced bivalence). This explains why they are likely to see man 924 as bald if they have just declared that man 923 is bald.

Nevertheless, since the participants in the exercise are competent speakers of English, we can be sure that they will not blindly go through the entire series, and call #2000, Jerry Garcia, bald. That way lies madness (or at least incompetence). Eventually, a few of the participants will demur from calling one of the men bald. If this group finds themselves in a small minority, they will go along with the majority, to maintain the consensus—for a while. As the group proceeds through the borderline area, more and more of the participants will demur, or feel like demurring. At some point, enough of them will demur that the consensus on 'this man is bald' will break down. Given the 'forced march' instruction, where they must achieve consensus on each answer, the group will eventually agree to call one of the men 'not bald'. Suppose this first happens with #975. Recall that they have just agreed, perhaps reluctantly for many of them, that #974 is bald, and so they put 'Man 974 is bald' on the conversational

score. Of course, when they then said that #975 is not bald, they did not contradict themselves. Nor have they violated the tolerance principle. At that point, the tolerance applies in reverse—we take the contrapositive. In declaring man 975 to be not bald, they implicitly deny that man 974 is bald, and so ‘Man 974 is bald’ is *removed* from the conversational record. It is similar to what happens when any presupposition is challenged (or contradicted) in the course of a conversation.

The event as described here is quite similar to the outcome in one of Lewis’s scenarios, recounted in the previous section. In that story, the participants in a conversation first agree to a ‘low’ standard when they accept ‘France is hexagonal’. Later, when they demur from ‘Italy is boot shaped’, the standard is raised, and so ‘France is hexagonal’ is implicitly removed from the record. Similarly, when the present conversationalists explicitly declare that #975 is not bald, they implicitly retract the statement that #974 is bald. In short, the conversational score is the device that enforces tolerance.

Just as ‘Man 974 is bald’ comes off the record, so does ‘Man 973 is bald’; ditto for a few more of their recent pronouncements. We cannot and thankfully need not set a precise boundary as to how many sentences are removed from the record once they jump. Exactly what is and what is not on the conversational record is itself a vague matter. Typically, borderline cases of ‘what is on the record’ do not interfere with a conversation. If a question arises about a specific case—say #967—the participants can ask each other about that fellow, and so the status of that case can be made explicit.

So let us continue the thought experiment by reversing the order of query. We assume that man #975 is the first ‘jump’, where our participants shift and deny that the man is bald. Suppose that we explicitly ask them about #974 again, after reminding them that they just called that man ‘bald’, and that they can barely distinguish #974 from #975 (if at all). Although I don’t intend to apply for a grant to confirm this empirical hypothesis, I’d speculate that they would explicitly retract that judgement, saying that #974 is not bald (and thus put ‘Man 974 is not bald’ on the record). The new consensus on the non-baldness of #975 will spread backward to cover #974. Suppose that we then ask them about #973. They



would retract that judgement as well. Then we can ask about #972. Of course, they will not move all the way back down the series, and end up denying that #1, Yul Brynner is bald. After all, he has no hair at all. At some point, they will jump again, and declare a certain fellow to be bald—suppose it is man 864. This again will result in the removal of certain items from the conversational record, such as the denial that man 865 is bald. If our conversationalists do not lose patience with us, we can then go back up the series, but the pattern is clear enough. They would declare men they encounter to be bald for a while, eventually jumping. In general, our conversationalists will move backward and forward through the borderline area. Tolerance is enforced at every stage, by removing judgements from the conversational record whenever a jump occurs.

Notice, incidentally, that pending empirical research, there is no reason to think that the participants will always jump at the same place(s) as they move back and forth through the middle part of the series. We only know that come hell or high water, they will eventually jump, well before they call Jerry Garcia bald and well before they deny that Yul Brynner is bald—assuming of course that they remain competent speakers of English and do not go batty as a result of this experiment. We know this, since we are ourselves competent speakers of English.

So how, exactly, is the sorites paradox resolved? Consider the version that uses a single, inductive premise:

for each  $i < n$ , if  $Pa_i$  then  $Pa_{i+1}$ .

It is outright false that for each  $i < 2000$ , if our conversationalists judge man  $i$  to be bald, then they will in fact judge man  $i+1$  to be bald. As they go through the forced-march, sooner or later they will jump. And this jump does not undermine their competence as speakers of English. Indeed, if they did not jump, and went on to call man 2000, Jerry Garcia, bald, they would thereby display incompetence. Moreover, as we have seen, a jump does not violate tolerance, since it comes with a retraction of items from the conversational record.

Something similar happens with the version of sorites that avoids a single, inductive premise in favor of each instance of it. Whenever the conversationalists jump from ‘bald’ to ‘not bald’, they thereby undermine one of the premises in the long argument.

Although the inductive premise is never true (in full splendor), notice that in the forced march version of the sorites, the conversational score never contains what may be called a ‘strong counterexample’ to the inductive premise—so long as tolerance is in force. That is, there is never a number  $n$ , such that ‘man  $n$  is bald’ and ‘man  $n+1$  is not bald’ are both on the record at the same time. If they jump to ‘not bald’ at  $n+1$ , then ‘man  $n$  is bald’ is retracted, and removed from the score. In the scenario as envisioned, tolerance is enforced on the conversational record.

This may just be a feature of the forced march, however. Suppose instead that the participants consider the men in the series in random order. Then they very well might judge man  $n$  to be bald and sometime later judge man  $n+1$  to be not bald without explicitly or implicitly retracting the first judgement. Indeed, they might even judge man  $n$  to be not bald and later judge man  $n+1$  to be bald—violating a penumbral connection. In fact, if they are not careful, they might even judge a man to be bald and a bit later judge the very same man to be not bald. In such cases, they have not noticed the violation of tolerance, penumbral connection, or consistency. But once the violation is pointed out, they will retract at least one of the judgements, and so one of the offending propositions will go off the record. One nice feature of the forced march is that tolerance, penumbral connection, and consistency are easily enforced. Not so in general.

In the forced march, there will surely be cases  $m$ , where, for example, ‘man  $m$  is bald’ is on the record and ‘man  $m+1$  is bald’ is not on the record. But that does not undermine tolerance, as that notion is defined above (§2). Since man  $m$  and man  $m+1$  differ only marginally, the participants cannot judge

them differently at the same time, but there is nothing to prevent them from judging one and not the other.<sup>2</sup>

It follows, incidentally, that the conversationalists cannot simultaneously judge every member of the series and remain competent in the use of the term ‘bald’, assuming again that tolerance remains in force. That is, the participants cannot have a judgement for every man on the conversational record all at once. If they did, they would either violate tolerance, or else they would call Yul Brynner not bald, or call Jerry Garcia bald. At any time, some of the baldness states will have to be left unjudged. That’s life. Perhaps they can get to a point where they have done a maximal amount of judging. Suppose, for example, that they judge every man from #1 to #974 to be bald, and they judge every man from #979 to #2000 to be not bald. Moreover, suppose that they cannot make a judgement about any of #975 through #978 without jumping—without retracting some of their previous judgements. Such is the nature of the series.

**5 Raffman on vagueness.** The present resolution of the sorites paradox has the same structure as that of Raffman (1994), (1996), and I acknowledge a deep debt to that penetrating work. Raffman envisions a single subject who is marched through a sorites series consisting of color patches. She is explicit, however, that her account is to apply to vague terms generally. To adapt her program to our series above, the subject would begin by calling the first few fellows bald, and continue that way a bit into (what I call) the borderline area. But she eventually jumps and calls one of the men not bald. Again, suppose that this first occurs at #975. Raffman says that at that point, the subject has changed to a different *psychological state*. In the new state, it is not true that, say, #974 is bald.

---

<sup>2</sup>As noted in §2 above, it would violate tolerance if the conversationalists judge man  $n$  to be bald, and *decide* to leave man  $n+1$  unjudged.

Once again, it is a commonplace that many properties are relative to context. For example, a lightweight internal lineman among high school football players would be an extremely heavy jockey. That is what Raffman calls an ‘external’ relativity to a comparison class. For vague predicates generally, Raffman proposes that there is an additional, internal relativity, to a psychological state. Patches look red, or are red only relative to the state of a competent person judging them to look red. A man is bald, or not, only relative to the state of a competent person judging whether he is bald.

Raffman argues that (something like the tolerance principle entails that) a competent subject cannot give different verdicts to ‘man 974 is bald’ and ‘man 975 is bald’ *at the same time*—and while in the same state. As above, if the subject is presented with any adjacent pair and asked to judge them both, she will agree that they are either both bald or neither bald: ‘at no time are adjacent [members] of the series simultaneously category-different’ (1994, 53). So it is never the case that an instance of the inductive premise is false. So long as the subject remains in the same state, she will not differ in her judgement concerning any adjacent pair of men in the series.

Drawing on some psychological research, Raffman (1994, 50) proposes that when the subject jumps during the forced march, she moves into a new psychological state. There is a ‘category shift’ in which the judged category—‘not bald’ in this case—‘spreads backward’ along a string of the preceding men in the series. In the new state, the extension of non-bald ‘expands backward, instantaneously’, to include some of the men that formerly fell in the extension of ‘bald’.

To adapt Raffman’s narrative to the present example:

My hypothesis is that, at the moment of judging [#975], the speaker undergoes a kind of *Gestalt shift* that embraces [#974] (and probably some of [his] predecessors) as well as [#975]. At the moment of shift to [‘not bald’], the speaker is disposed to judge both [#974] and [#975] (plus some [men] on both sides) as being [not bald], thereby allowing for a change in kind while preserving the effective continuity of the series. Intuitively speaking, a string of [men] shift their

[baldness state] simultaneously, so that [#974] and [#975] never differ . . . *at the same time*. Like the duck-rabbit and Necker cube, these [men] can ‘go either way’: they can be ‘seen as [bald]’ or ‘seen as [not bald]’—now one way, now the other . . . If asked to reverse direction and retrace his steps down the series toward #1, the speaker would now judge [as not bald] some [men] that he previously judged [bald]. At some point, of course, he would shift back to [‘bald’]; for example he might judge [#974] through [#903] [not bald], but then undergo a Gestalt switch back to [‘bald’] at [#902]. And so forth . . . (Raffman 1996, 178)

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. (Raffman 1996, 180)

The reference to Gestalt shifts is suggestive. A subject’s perception of a Necker cube, for example, changes without any changes in the drawing itself. With a vague predicate, the judgement shifts with a small but perhaps imperceptible change in the item being judged.<sup>3</sup>

The relativity to psychological state allows for a distinction between a ‘categorical judgement’, which consists of a subject considering a single case, and a ‘comparative judgment’, where the subject deals with two (or more) instances at once. Raffman illustrates this, and the relativity, by envisioning a forced march in which the items are presented in pairs. Let us return to our series of 2,000 men. We first ask the subject, or the participants, to judge the pair <#1,#2>. Both are judged to be bald. Then we present the pair <#2,#3>. Both bald. Then <#3,#4>, and so on. Since the pairs are marginally different, perhaps indistinguishable, the subject(s) will always judge each pair alike. But, again, the subject(s) will eventually jump. They won’t go on to call Jerry Garcia and his predecessor both bald. Suppose that the pair <#887,#888> are judged to be bald, but the pair <#888,#889> are judged to be not bald a few

---

<sup>3</sup>Thanks to an anonymous referee.

seconds later. If we do not allow for a contextual shift of some sort—be it a change in psychological state, a change in the conversational record, or something else—the subjects will have contradicted themselves at that point. Since man 888 is in both pairs, he is judged to be bald and then judged to be not bald a moment later.

There are not many options for interpreting this scenario. The theorist can conclude that vague predicates like ‘bald’ have some contradictory cases, or that any language (like English) with vague predicates is incoherent generally (following Dummett 1975). Or the theorist can find fault with the scenario. Perhaps there is something illegitimate in demanding a forced march through a sorites series. I presume that Williamson (1994) would counsel this reply. According to this staunch epistemicist, the subject(s) should refuse to answer when they get near the (precise) border, despite their instructions. Since they do not and cannot know where the border is, they have no business venturing an opinion on the state of baldness of the fellows near that border. Or the theorist can claim that the subject(s) are incompetent, or have made some sort of mistake with their answers. But what sort of mistake? The subjects are told to call them as they see them. What did they do wrong? Or the theorist can follow Raffman and the present program and say that a shift of context has occurred. Man 988 is bald in one context and not bald in another.

Raffman’s contextualist framework complements the present one. She sees the extensions of vague predicates as varying with the psychological states of (competent) speakers of the language, while I see the extensions as varying with conversational records among (competent) speakers. In one sense, her thought experiment is a special case of the present one. We might think of hers as a limiting case of a ‘conversation’ consisting of only one participant. The ‘conversational score’ would contain features of the psychological state of the subject.

Of course, a conversation consists of its individual members, and Raffman’s notion of psychological state is intended to explain what happens in the mind of those individuals. The communal,

conversational analogue of Raffman’s gestalt-shift is what I call a ‘jump’, the combination of a breakdown of consensus among the members of the conversation, and the forging of a new consensus on a new verdict. A consensus breaks down only if enough of the individual participants demur from the communal judgement, and if Raffman is correct this happens when enough of them have made the relevant gestalt-shift (and refuse to switch back even for the sake of conversational harmony). So if Raffman’s account can be sustained, it is more basic than mine. I have little to contribute to the articulation of the relevant notion of psychological state, nor to the psychological explanation of individuals as they judge vague predicates.

The present account focuses instead of the communal aspects of language use, such as the role of language in public communication. Surely, it is individuals who communicate, and these individuals are in different psychological states at different times. I do not have much to say by way of a detailed account of how this works. What is the bridge between the different psychological states of individual communicators and the public side of communication? The cooperative features of conversations that Lewis (1979) sketches and, in particular, the defeasible drive for consensus, feed into the role of conversational record in handling borderline cases of vague predicates. They suggest the more communal sort of relativity that the present work focuses upon.

**6 Determinacy revisited, truth, and classical semantics.** The time has come to further articulate, or refine, the notion of determinacy, to indicate the contextual elements that are relevant to its extension.<sup>4</sup> Let  $F$  be a monadic (possibly complex) predicate in a natural language. Recall the previous definition of determinacy, quoted from McGee and McLaughlin (1994, §2):

---

<sup>4</sup>I am indebted to Agustin Rayo here.

. . . to say that an object  $a$  is [determinately] an  $F$  means that the thoughts and practices of speakers of the language determine conditions of application for . . .  $F$ , and the facts about  $a$  determine that these conditions are met.

A first and hasty interpretation of this is that  $a$  is determinately an  $F$  if and only if the (linguistic) meaning of the predicate  $F$  and the non-linguistic facts determine or guarantee that  $Fa$  holds. In short,  $Fa$  is determinate in this sense if its truth supervenes on meaning and non-linguistic fact. Since English is a public language, it seems plausible that the meaning of its words is at least largely independent of the context of utterance (ambiguity and considerations like those in Davidson 1986 notwithstanding). Otherwise, we are left with the specter of Humpty Dumpty's claim to Alice, and no real communication. And, of course, most non-linguistic facts (such as the amount and arrangement of hair on a person's head) are independent of the context of utterance. So the present interpretation of determinacy would allow virtually *no* contextual elements into the notion. Suppose, for example, that Harry just went to see a baseball game, the only game in town at the time. No vagueness about that (or none that matters here). Still, the sentence 'Harry went to the game' would not be determinately true, since nothing about the meaning of the words and the relevant facts about Harry, the players, etc., determines that this sentence is true. The (linguistic) meaning of the phrase 'went to the game' does not fix which game we are talking about. Similarly, on the present interpretation, the sentence 'Everyone is present and accounted for' would not be determinately true (in any context), since the meanings of the words by themselves do not fix who we are talking about. The sentence, 'I am hungry', would also fail to be determinately true in any context, since meaning and relevant facts do not determine the referent of 'I'.

McGee and McLaughlin did not have such an austere notion of determinacy in mind. Surely, 'the thoughts and practices of speakers of the language' conspire to fix some items in various contexts of utterance. In the first example, the members of the conversation would determine which game they are talking about, by what they say, do, and think. So 'Harry went to the game' would be determinately true,



on the McGee and McLaughlin notion. Similarly, in a given context, the sentence ‘everyone is present and accounted for’ would be determinately true or determinately false (unless there is some vagueness involved). The conversation would fix the extent of ‘everyone’, and this information would be on the conversational score. If a hungry speaker says ‘I am hungry’, then what she says is determinately true.

I conclude, then, that at least some contextual elements should be included in the analysis of determinacy. In particular, when it comes to vagueness, determinacy is sensitive to the comparison class, the paradigms, or the comparison cases. Suppose, for example, that during a conversation about NBA players, somebody says that a player who is six feet, one inch tall is short. That sentence is determinately true (in that context). The conversational record would indicate the comparison class, and given that class, the sentence is true.

So far, I presume that everything is in line with the notion of determinacy invoked by McGee and McLaughlin (1994). However, I suspect that I am about to part company with them. I propose that the borderline cases of a vague predicate that have been decided in the course of a conversation *not* be included in what fixes determinate truth. Suppose that the relevant comparison class, paradigm cases, or comparison cases for ‘baldness’ are fixed. Assume that Harry remains a borderline case. According to the foregoing open-texture thesis, in some situations, speakers are free to assert ‘Harry is bald’ without undermining their competence, and they may likewise assert ‘Harry is not bald’ in some situations. Unsettled entails open. So suppose that someone asserts ‘Harry is bald’, in the course of a conversation, and this assertion goes unchallenged. As above, the sentence (or proposition) ‘Harry is bald’ goes on the conversational score. One might take it that ‘Harry is bald’ would then be *determinately* true in the context after the utterance, at least until the assertion is retracted (in which case the context changes). After all, once the utterance is accepted onto the record, the thoughts and practices of (those) speakers have fixed a truth value for this sentence. Nevertheless, I propose that ‘Harry is bald’ is *not* determinately true in the context of this conversation.

As noted in §5 above, Raffman (1996) distinguishes the ‘external’ context of a judgement (concerning a vague predicate), which fixes the comparison class and the like, from the internal context, consisting of the subject’s psychological state. Just about everyone accepts that the extensions of vague predicates can vary with the external context. Raffman’s thesis is that these extensions also vary with the internal, psychological context. Once again, I invoke a similar distinction here, in the broader perspective of public conversations. The present proposal is that ‘external’ factors—comparison class, paradigm cases, comparison cases, etc.—are included in what fixes determinate truth and determinate falsehood, but internal factors, and in particular, the decided borderlines cases, do not modify the extension of determinate truth and determinate falsehood.

When it matters, I will use the phrase ‘e-determinate’ to indicate the present notion, and distinguish it from the broader McGee-McLaughlin notion, where *all* contextually fixed factors, including decided borderline cases, are included. To mark the contrast, I’ll use the word ‘established’ for the broader McGee McLaughlin notion.

Since the word ‘determinacy’ is a term of art, I am free to define it as I please, so long as the definition is coherent (even if vague). Nevertheless, the reader deserves a word of justification for my notion of e-determinateness. One advantage of the broader notion of ‘established’ is that it allows, or may allow, the imposition of ordinary classical semantics. That is, within a fixed context, one can coherently maintain that vague predicates have precise extensions, and complementary anti-extensions. Raffman, for example, uses counterfactual conditionals to fix the extensions. Consider, for example, our paradigm sorites series, consisting of 2,000 men ranging from Yul Brynner to Jerry Garcia. Fix what Raffman calls the external context, and consider a single competent subject in a particular psychological state *s*. Then a given man *m* is in the extension of ‘bald’ in *s* if it is the case that the subject would judge *m* to be bald when in state *s*, and *m* is in the anti-extension of ‘bald’ otherwise. To be sure, at any given time, neither the observing subject nor an outside examiner can learn the baldness statuses of all of the

men relative to the psychological state  $s$ , or at least not in the straightforward manner. As noted, for some of the men in the series, if the subject actually did judge them to be bald, or not bald, she would no longer be in the state  $s$ . Some judgements trigger a shift to a new psychological state. Once that happens, we are no longer able to determine the status of the other men in the original state  $s$  (unless we could get the subject back into  $s$ ). Nevertheless, on Raffman's view, for each external context and each psychological state, there is a fixed and precise extension for each vague predicate in that state, fixed by the indicated conditionals, most of which are counterfactual.

Raffman's proposal might be imported full blown into the present conversational context. Let us define a man in our series to be in the extension of 'bald' at a given stage in a conversation, if the conversationalists would judge him to be bald, if they were asked that question. In the limit, however, we can actually determine the status of only a few of the men that way. Once we ask the conversationalists about any of the men in the borderline area, and get an answer, we thereby change the context, since the conversational record changes at that point. We are now at a new stage in the conversation, and all bets are off concerning the original stage. Nevertheless, the status of each man in the series is fixed at any given time by how the subjects *would* respond, if queried about that man.

Another prominent contextualist, Delia Graff (2000), holds that at any given time, there are sharp borders in the sorites series. The problem is that the borders shift with the focus of the subjects. The reason a subject never sees or notices the border is that it is never located at the place where she is looking—the border is never (and can never be) salient. Such is the nature of contextually determined vague predicates. Graff's view is an epistemicism of sorts.

With the imposition of precise extensions (and complementary anti-extensions) in each context, the foregoing contextualists thus invoke an ordinary, classical model-theoretic semantics. With that comes classical logic, a neat package.

To accomplish this, however, the Raffman framework assumes that the counterfactuals are all well-defined, and each has a unique truth value. In particular, there must be a fact of the matter concerning how the subject would respond in the given psychological state, if queried about a given case. If there is any vagueness, or indeterminacy, in the counterfactuals, the imposition of classical semantics might fail. The same goes for the potentially similar treatment in the conversational context. There must be a fixed fact of the matter concerning how the conversationalists would respond to each query. Similarly, Graff's imposition of classical semantics stands or falls with the assumption that at any given time, there are indeed precise, but unknowable borders that lie outside of the range of consideration.

I do not intend to challenge these assumptions or premises of my fellow contextualists. Let me just register skeptical agnosticism concerning the presuppositions, and then concede, for the sake of argument, that we can indeed achieve classical semantics, and classical logic, if we invoke the broad notion of determinacy (i.e., establishment). That is, I assume for the sake of argument that if the external context and the internal psychological state are both held fixed, then we can and should reason classically.

Note, however, that the relevant internal context can change very rapidly. In Raffman's case, it is an empirical question exactly when, and how often, a given subject's psychological state changes. At a minimum, the state changes whenever one of the counterfactuals:

if asked about man #*i*, the subject would respond 'bald'

changes its truth value. It is not implausible that this happens every moment, or every few moments (given how delicate things can get in the borderline area, see §8 below). In the present conversational framework, the context changes every time something is added to the score, which happens every time something not already established is said. For Graff, the extension of vague predicates changes every time the subject's attention shifts more than a little.

This is not a complaint against classical logic. It is a commonplace that one can reason classically only so long as the references of the terms and the extensions of the predicates in an argument remain fixed. We remind introductory students of this, noting the ancient fallacy of four terms and the like. On the views under discussion, the extensions change with every change in internal or external context, and we can reason correctly only if we are aware of this.

This observation also explains why the sorites reasoning seems so tempting, and it indicates how to resist the key premise. However, the constant shifting of context puts a damper on the appropriateness of classical logic when reasoning with vague predicates (if indeed it is appropriate). The supposedly precise extensions of the predicates can change in the very act of our considering them as we go through an argument, trying to reason in an ordinary context. We might call this a Heraclitus problem. The river changes every time we step into it. The extensions of vague predicates change momentarily, right before our eyes, even as we are considering the premises of an argument.

As noted, I will not mount a challenge to the assumptions that lead to the imposition of classical model-theoretic semantics and classical logic. Nevertheless, I am interested in a notion of ‘context’ and of correct reasoning therein, where the shifts are not so erratic. It seems to me that it is plausible to hold the *external* context fixed for a period of time, since that does not change quite so rapidly, nor need it change without notice. Again, the external context fixes the comparison class, the paradigm cases, and/or the comparison cases (as well as the extensions of proper names and pronouns like ‘everyone’). I submit that this is the normal background of reasoning with vague predicates, and I am interested in the norms of correct reasoning in such contexts. Let the actual extensions of vague predicates vary within such a context, as we reason about them. For this purpose, the current notion of e-determinacy is an interesting and important one for the logic of vague predicates. Unfortunately, this can be pursued in detail only on another occasion.

It is clear that e-determinate truth is not the same thing as truth. There is no analogue of the supervaluationist slogan that truth is super-truth. Suppose, again, that external factors are fixed and that Harry remains a borderline case of baldness. And suppose that a given competent subject judges Harry to be bald, in a given conversational context. Then it is true *in that context* that Harry is bald. Nevertheless, it is not e-determinately true that Harry is bald in that context. As far as e-determinacy goes, Harry remains a borderline case.

There is, I presume, a norm of conversation that one should only assert truths (other things equal, of course).<sup>5</sup> The norm is *not* that one should only assert e-determinate truths. That would preclude us from deciding any borderline cases, even temporarily. The norm is that one should only assert truths, and on the present account, truth varies with context. If Harry is borderline bald, and nothing on the conversational score (thus far) precludes judging Harry to be bald (e.g., it is not the case that ‘Harry is not bald’ is on the score), then a speaker cannot help satisfying the norm in sincerely asserting (or denying) Harry’s baldness. The acceptance of ‘Harry is bald’ onto the conversational score thereby makes the sentence true in that context, and thus automatically satisfies the norm. Of course, this holds only for items in the borderline area. Anyone who calls Jerry Garcia bald or denies that Yul Brynner is bald has violated the norm.

**7 What does all this conversational stuff have to do with semantics and logic?** A common, perhaps natural, response to the present program (and to Raffman’s) is that it focuses on the pragmatics (or psychology) of the use of vague terms, and has nothing to do with semantics, and thus nothing to do

---

<sup>5</sup>I am indebted to Patrick Greenough here.

with logic.<sup>6</sup> The underlying objection seems to be that the correct resolution of the sorites paradox should turn exclusively on the meaning of the terms. Logic flows from meaning.

This presents a false dilemma. Raffman (1994, 43) writes that her ‘story is at bottom a psychological one, resting on a hypothesis about the mental representations that underlie our usage of vague words’ and so ‘where vague predicates are concerned, logic and semantics are more intimately entwined with psychology than might have otherwise been supposed’. The present story turns more on the pragmatics of conversations involving vague terms, and so I would conclude instead (or, better, in addition) that where vague predicates are concerned, logic and semantics are intimately entwined with pragmatics. Along similar lines (perhaps), Wright (1987, 277) proposes that we seek a semantics according to which linguistic competence is understood ‘on the model of a practical skill, comparable to the ability to . . . ride a bicycle’. Of course, unlike bicycle riding, we exhibit the practical skill of language mastery in concert with each other (more or less). Otherwise, there would be no communication which is, after all, the point of this enterprise. The meanings of vague terms are intimately tied up with the proper display of this skill.

As Humpty Dumpty ought to have put it, competent users of the language are its masters. I submit that the correct use of vague terms is bound up with psychology and pragmatics. The notion of open-texture (see §2 above), which is shared between the present account and Raffman’s (and Wright’s), shows how this is so. This requires at least some revisions of some widely-held views on semantics, meaning, and extensions.

Raffman (1994, 69-70) articulates the following biconditional:

---

<sup>6</sup>Much of the argument in this section must be heavily revised, due in large part to the criticisms brought in Rosanna Keefe’s contribution to this volume, and due to conversations with Diana Raffman and Crispin Wright.

(B) [A]n item lies in a given category if and only if the relevant competent subject(s) would judge it to lie in that category.

See also Wright (1976), (1987).

Although Raffman only invokes (B) in the context of vagueness, I would think that for a vast range of predicates—vague or not—in public languages, the thesis is in the neighborhood of a truism. A piece of metal is gold if and only if relevant competent subjects would judge it to be gold (if asked); an animal is a marsupial if and only if relevant competent subjects would judge it to be a marsupial; a natural number is prime if and only if relevant competent subjects would judge it to be prime. And a man is bald if and only if relevant competent subjects would judge him to be bald.

My suggestion that many cases of (B) are true depends crucially on what it is to be a ‘relevant competent subject’. The claim, so far, is only that in a lot of cases, it is possible to specify what it is to be a relevant competent subject in such a way that (B) comes out largely true. Moreover, the ‘if and only if’ in (B) is only a material biconditional. No causal or semantic connections are postulated.

In each case, even after the class of relevant competent subjects is specified, so that (B) is largely true (when it is), there is a question as to which is the chicken and which the egg. It is an instance of what Wright (1992, 108-140) calls the *Euthyphro contrast*. To adapt one of Wright’s examples, consider the following instance of (B):

A story is funny if and only if competent subjects would judge it to be funny.

In this case, competent subjects are those with normal senses of humor, and I presume that with this stipulation, there is no controversy over the truth of this instance of (B). It is plausible that in this case, the judgements of these competent judges are somehow *constitutive* of funniness. What makes a story funny is that people tend to find it funny. Plausibly, then, humor is response-dependent, or, to paraphrase Raffman, judgement-dependent. Let us call this the *Euthyphro* reading of (B).

Contrast this with another case of (B):



A number is prime if and only if competent subjects would judge it to be prime.

In this case, the relevant competent subjects are those good at arithmetic, or perhaps (some) mathematicians. With this stipulation, the material biconditional is largely true, at least for smallish numbers. But in this case, of course, it is not the judgements of these competent subjects that *make* the number prime.<sup>7</sup> A number is prime if it has exactly two (distinct) factors: itself and 1. A subject is competent here only if she gets it right most of the time, other things equal. In this case, competent judges *track* the truth, their judgements do not constitute it. Call this the *Socrates* reading of (B).

A Euthyphro reading of an instance of (B) is a thesis about the extension of the predicate in question ('funny'). A Socrates reading is a thesis about what it is to be competent. To be sure, I do not claim that this distinction is a sharp one, nor that it is an all or nothing matter. I presume that many predicates have a mixture of Socratic and Euthyphronic factors, and perhaps there are no, or very few. pure cases of either.

OK, what of our current example?:

A man is bald if and only if competent subjects would judge him to be bald.

In this case, the specification of 'competent' is rather broad. A subject is competent if she understands the language, and is accurately perceiving the man in question under normal conditions (so she can see at least roughly how much hair each of them has and how it is arranged). Raffman (1994, 70) writes that thesis '*B* is true with respect to borderline cases because our competent judgements of borderline cases determine their category memberships; conversely *B* is true with respect to clear cases because clear cases determine what counts as competent judgement.' In present terminology, if the predicate is vague,

---

<sup>7</sup>What is at stake here is the objectivity of the relevant discourse (see Wright 1992). Readers sympathetic to (semantic) anti-realist accounts of mathematics should pick another example. The present distinction is more complicated, if available at all, for anti-realists about all discourse.

then (B) is at least largely Socratic for clear cases. If someone asserts that Yul Brynner is not bald, or that Jerry Garcia in his prime is bald, then we would regard her as at least temporarily or locally incompetent—either as not understanding the meaning of ‘bald’, or not looking carefully, or misperceiving. But when it comes to borderline cases, (B) is to be read with Euthyphro. Since, by open-texture, borderline cases can go either way, the judgements of otherwise competent subjects *determine* whether the man is bald—in the relevant conversational (or psychological) context.<sup>8</sup> As Raffman (1994, 44) puts it early on, ‘an adequate treatment of vague predicates and their sorites puzzles must appeal to the character of our judgements about the items in the series’. Every vague predicate is judgement-dependent *in its borderline area*.

In any case—and on either reading—thesis (B) is not a statement about the *meaning* of the predicate in question (whether it is vague or not). Raffman (1994, 58) writes that the ‘sorites . . . is solved independently of any particular meaning analysis of the predicate . . . On the contrary, all that is required to solve the puzzle is a claim about the correct application or *extension* . . . of the predicate at issue.’ A bit later, she elaborates:

I do not claim that the *meaning analysis* or *intension* of a vague predicate includes a judgmental element. For instance, I do not claim that in calling an object red one means or is saying, in either the ‘speaker’ or ‘semantic’ sense, that the object is merely red-relative-to-me-now or red-relative-to-such-and-such-a-context . . . Rather, I claim that the *extension* of ‘red’—the class of objects that satisfy the predicate—is always relativized to certain psychological (and nonpsychological) contexts. The sorites is a puzzle about the correct *application* of vague predicates, and that is all my story addresses. (Raffman 1994, 66)

---

<sup>8</sup>To be sure, ‘borderline bald’ and ‘competent speaker’ are both vague. Thus, it may not be determinate in a given case whether (B) is to be read with Socrates or with Euthyphro.

The first entry in *Random House Webster's unabridged dictionary* for 'bald' reads 'having little or no hair on the scalp'.<sup>9</sup> There is no reason to challenge this definition, which makes no mention of—or even implicit reference to—the judgements of competent users. Clearly, however, what *counts* as having little hair on the scalp depends on what Raffman calls the 'external' context. Is it a sales meeting of a company that makes shampoo, or a sales meeting of a company that makes toupees, or a sales meeting of a company that makes sun-screen? And if the foregoing account is correct, what counts as having little hair can also vary with the shifting *internal* context of a conversation (or, on Raffman's account, it varies with psychological state).

When Alice objected to Humpty Dumpty's claim that his words mean just what he chooses them to mean, he replied that it is a question of who 'is to be master'. Surely, our sympathies here are with Alice. Individual speakers are not 'masters' of the *meanings* (or intensions) of the words they use, independently of the thoughts and actions of the wider community of language users (up to the compelling points in Davidson 1986). But in the borderline region, individual, competent speakers are indeed 'masters' of the *extensions* of vague terms. Humpty Dumpty is right in this limited domain.

**8 Underground in the borderline area.** Recall the sorites series introduced in §4 above, which consists of 2000 men ranging from Yul Brynner to Jerry Garcia in his prime. We envisioned a forced marched conversation through this series, starting with Yul Brynner. Given their instructions, and assuming their competence in English (and normal observation), the participants to the conversation will eventually jump and call one of the men not bald. Again, suppose that this first happens with man 975. If we were to run the same series with the same group of English speakers (in the same order) on another

---

<sup>9</sup>The relevant entry in the *Oxford English dictionary* reads in part 'having no hair on some part of the head where it would naturally grow'.

occasion, they might jump at a different place, say #967 or #984. They would almost certainly jump at a different place if they were marched through the series in the reverse order, starting with Jerry Garcia, and moving down.

Presumably, something causes the jumps when they occur. The events are too macroscopic for quantum randomness to be involved (or are they?). The point here is that so long as the man being judged for baldness is in the borderline area, the meanings of the word ‘bald’, the external contextual factors, and the non-linguistic facts (about the number of and arrangements of his hair) do not determine a correct response. So as far as the normativity of meaning goes, they can go either way. Raffman (1996) emphasizes that the exact location of the jump carries ‘no normative force’; it underwrites ‘no distinction between correct and incorrect usage’. A different group, or the same group on another occasion, that jumps at a different place is not thereby in error. Raffman elaborates further:

... there is no reason to [jump], hence no justification for [jumping], 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.

By ‘reason’, Raffman presumably means a ‘semantic reason’. It is virtually analytic that if semantic rules (and non-linguistic facts) always determine when the jump is supposed to occur, then the predicate is not vague after all—there are no borderline cases.

The situation is (vaguely) reminiscent of the old saw about Buridan’s Ass. The animal finds itself exactly midway between two bales of hay. Since he does not have any reason to walk toward the first bale (as opposed to the second), he does not approach that bale. Similarly, he does not have any reason to go to the second bale (as opposed to the first). So the poor animal remains where it is, and starves to death. An unfortunate and irrational outcome. He really should have gone one way or the other, and it does not matter which. Similarly, our conversationalists do not have a semantically

compelling reason to jump at man 974, nor do they have a semantically compelling reason to jump at #975, nor do they have a semantically compelling reason to jump at #976, etc. But they are required to jump *somewhere*, and they become semantically incompetent if they do not.

Raffman (1994, 46) writes:

Now one thing we know is that at some point on each (complete) run of judgements along the series . . . a [jump] just does occur: the subject's slide down the slippery slope is broken . . . Just where the shift does occur on any given run will depend on a constellation of factors, including the direction in which the subject proceeds along the series, where in the series he begins his judgements, his perceptual state at the time, and so forth . . . But occur it does, on every run.

It is straightforward to adapt these observations to the present, social conversational context.

In a given case, the exact location of the jump is a 'brute mechanical' matter (Raffman 1994, 65). Presumably, it will depend on how tired or attentive the subjects are, what else is on their minds at the time, their emotional state, and countless other factors. To be sure, these factors have nothing to do with the amount and arrangement of hair on the heads of the men in the series, and so are irrelevant to the meaning and thus the proper application of 'bald'. But they are most relevant to its *extension* in the given context. They are relevant to what *counts as* having little or no hair on the scalp, in the given context.

Raffman (1994, 53) adds that what 'will or will not trigger the [jump] is not something to which we, as judging subjects, have access; so far as the subject is concerned, the [jump] simply occurs'. Presumably, the members of the conversation can feel the shift coming as they march through the sorites series. When they move through the borderline area, they will find themselves less and less comfortable with their judgements that the men are bald, and the group will find it harder and harder to maintain consensus on baldness. The jump is not a blind process, in the sense that they have no idea when it will occur. After all, they are in charge of the situation—masters as Humpty Dumpty would put it. But the

jump is blind in the sense that the subjects need not be aware of the factors that actually trigger the shift—i.e., when enough is enough.

An adaption of an example from Wright (1987) illustrates the situation. Consider a digital tachometer connected to a motor. This is a ‘brute mechanical’ device if anything is. Suppose the device registers in units of 10 rpm, and suppose that the motor is running smoothly 300 rpm. The tachometer dutifully reads ‘300’. Now we slowly increase the speed of the motor in increments of 0.1 rpm. Of course, the tachometer will jump to 310 somewhere. Say it happens when the motor hits 306.2 rpm. If we speed the motor up a bit more, and then start slowing it down, the tachometer will jump back to 300 somewhere. But there is no reason to think that this jump back will occur at 306.2. In a typical case, the tachometer will keep its holding at 310 for a bit, and not jump back to 300 until later, say 305.7 rpm. If we did the experiment on another occasion—the next season, say—we might get slightly different results. It might first jump to 310 at 306.4 rpm and it might jump back to 300 at 305.8.

On each occasion, the exact location of the tachometer’s jump from 300 to 310, or back, depends on a number of factors, such as the ambient temperature, relative humidity, wear and tear on the parts of the tachometer, fluctuations in the power source, etc. Of course, these things have nothing to do with the speed of the motor, and so are irrelevant to the correctness of the reading. For the practical purposes at hand, this does not matter. So long as the seemingly random fluctuations affect the reading only in the ‘borderline’ area—say roughly 302.5 to 307.5 rpm—they will not keep the tachometer from operating correctly. When the motor is running at or near a ‘clear’ case, say 309.5 rpm, then these fluctuations do not affect the reading.

The analogy with our conversationalists on a forced march is straightforward. When they are in the borderline area, they can go either way without impugning their competence, and which way they go in that area depends on a number of factors irrelevant to the baldness state of the men in the series.

When they are considering a man who is not in the borderline area, then these side factors do not affect them, and they judge correctly (if they are competent and observing properly).

With the tachometer and the sorites series for baldness (or redness), the exact location of the jumps does seem to be completely arbitrary, in that it really does not matter *at all* where the jumps occur (so long as they occur in the borderline area). With some sorites series, however, the jump point is not arbitrary. Consider a sorites that goes from a spatially separated sperm-egg pair to a resulting two year old baby developed from that pair, with increments of .01 seconds. The location of the borderline of the predicate ‘person’ or ‘human being’ (or ‘being with human rights’) has obvious moral ramifications. People who put the border in one place rather than another are not acting arbitrarily. They may well have compelling reasons for what they do. The point here is that there are no *linguistic* reasons for putting the border in one place rather than another. Presumably, both sides of the abortion debate are speaking the same language. If I consistently put the border in the region of viability, and a mob comes and burns down my house as a result, they will not be accusing me of linguistic incompetence (whatever else they may accuse me of).<sup>10</sup>

Returning to our paradigm cases, for both the conversation about baldness and the tachometer, the jump-point is a ‘brute mechanical’ process, and there is no principle of tolerance for it. Nature does allow sharp jumps (or at least apparent discontinuities). Raffman (1994, 56) writes that ‘we have no *intuition* that if [#874] does not trigger . . . a [jump] then neither will [#875]’, just as we have no intuition that if raising the motor to 306.1 rpm does not cause the tachometer to jump to 310, then neither will 306.2 rpm.

---

<sup>10</sup>I am indebted to Carl Posy here. Notice that in any case, there will still be some arbitrariness in the ‘jump’ point. Viability and, to a lesser extent, conception are themselves vague matters.

To be sure, if our conversationalists jump too close to the clear cases, their competence as speakers of English (or their powers of observation) is compromised. Similarly, if our tachometer jumps at the wrong place—say it reads 310 when the motor is running at 300.2 rpm—its accuracy will be compromised, and possibly also its usefulness (depending on what it is to be used for). Although there is no principle of tolerance in the operation of a mechanical device, the *usefulness* of such a device is a context-dependent, and probably also a judgement-dependent matter. How accurate the tachometer need be to be useful is a vague matter, and a principle of tolerance does hold for that.

So how close to the border is ‘too close’ for competent speakers to jump? How inaccurate can the tachometer be without being ‘too inaccurate’ for a given purpose? Since these are also vague matters, one might think that we have just postponed (or moved) the problem of this work. This raises a problem in the neighborhood of higher-order vagueness, which is a topic for another day.

ACKNOWLEDGMENTS. This paper is a pilot for a longer work, which includes a model-theoretic semantics and a treatment of so-called higher-order vagueness. My biggest debt is to my colleague and friend, Diana Raffman, whose work inspired the present account, and who continues to give me valuable advice. Thanks also to the members of a graduate seminar on vagueness that we gave together in the Spring of 2002: Julian Cole, Sven Walters, Steven James, Jack Arnold, and Michael Jaworski. I am also indebted to the audiences at the ‘Liars and Heaps’ conference held at the University of Connecticut in the Autumn of 2002, an Arché workshop and subsequent seminar at the University of St. Andrews, the Hebrew University Logic Colloquium, and the Philosophy Colloquium at the University of Maryland. I have received valuable feedback from Carl Posy, Mark Sainsbury, Crispin Wright, Agustin Rayo, Patrick Greenough, Delia Graff, Barbara Scholz, Tim Williamson, Neil Cooper, Roy Cook, Rosanna Keefe, Brian Weatherson, Graham Priest, Michael Morreau, and Georges Rey.



## REFERENCES

- Davidson, D. (1986), 'A nice derangement of epitaphs', in Richard Grandy and Richard Warner, editors, *Philosophical grounds of rationality*, Oxford, Oxford University Press, 157-174.
- Dummett, M. (1975), 'Wang's paradox', *Synthese* 30, 301-324; reprinted in Keefe and Smith Keefe and Smith (1997, 99-118).
- Fine, K. (1975), 'Vagueness, truth and logic', *Synthese* 30, 265-300; reprinted in Keefe and Smith (1979, 119-150).
- Graff, D. (2000), 'Shifting sands: an interest-relative theory of vagueness', *Philosophical Topics* 28, 45-81.
- Horgan, T. (1994), 'Robust vagueness and the forced-march sorites paradox', *Philosophical Perspectives* 8: *Logic and Language*, 159-188.
- Keefe, R. and P. Smith (1997), *Vagueness: a reader*, Cambridge, Massachusetts, MIT Press.
- Lewis, D. (1979), 'Scorekeeping in a language game', *Journal of Philosophical Logic* 8, 339-359.
- McGee, Vann (1991), *Truth, vagueness, and paradox*, Indianapolis, Hackett Publishing Company.
- McGee, Vann and B. McLaughlin (1994), 'Distinctions without a difference', *Southern Journal of Philosophy* 33 (supplement), 203-251.
- Priest, G. (1987), *In contradiction: a study of the transconsistent*, Dordrecht, Martinus Nijhoff Publishers.
- Raffman, D. (1994), 'Vagueness without paradox', *Philosophical Review* 103, 41-74.
- Raffman, D. (1996), 'Vagueness and context relativity', *Philosophical Studies* 81, 175-192.
- Sainsbury, R. M. (1990), 'Concepts without boundaries', Inaugural Lecture, published by the King's College London, Department of Philosophy; reprinted in Keefe and Smith (1997, 251-264).
- Unger, P. (1975), *Ignorance*, Oxford, Oxford University Press.

Williamson, T. (1994), *Vagueness*, London and New York, Routledge Publishing Company.

Wright, C. (1976), 'Language mastery and the sorites paradox', in G. Evans and J. McDowell, editors,

*Truth and meaning: essays in semantics*, Oxford, Oxford University Press, 223-247; reprinted

in Keefe and Smith (1997, 151-173).

Wright, C. (1987), 'Further reflections on the sorites paradox', *Philosophical Topics* 15, 227-290;

reprinted in Keefe and Smith (1997, 204-250).

Wright, C. (1992), *Truth and objectivity*, Cambridge, Massachusetts, Harvard University Press.