

16

Saying More (or Less) Than One Thing

Andrea Iacona

In a paper called *Definiteness and Knowability*, Tim Williamson addresses the question whether one must accept that vagueness is an epistemic phenomenon if one adopts classical logic and a disquotational principle for truth. Some have suggested that one must not, hence that classical logic and the disquotational principle may be preserved without endorsing epistemicism.¹ Williamson's paper, however, finds 'no plausible way of substantiating that possibility'. Its moral is that 'either classical logic fails, or the disquotational principle does, or vagueness is an epistemic phenomenon'.² The moral of this chapter, on the contrary, is that there is a plausible way of substantiating that possibility. The option it contemplates looks like a view that Williamson dismisses at the beginning of his paper, and that others regard as unworthy of serious consideration.



16.1

A couple of preliminary clarifications. The first concerns the expression 'borderline case'. It is widely accepted as a matter of definition that a vague word is a word that admits of borderline cases. Typically, a vague predicate is a predicate that admits of borderline cases. Yet this leaves unsettled what 'borderline case' means exactly. According to one reading of the expression, a borderline case is a *situation* that involves actual problems of evaluation. We are 'in' a borderline case so understood when a speaker assertively utters a sentence and we don't know whether the assertion is correct or incorrect. For example, it may happen that a speaker assertively utters 'A is tall' and we don't know how to evaluate the assertion, because we don't know whether it is correct to say that A falls into the extension of 'tall'.

According to another reading of the expression, a borderline case is an *object* that neither clearly belongs to the extension of a predicate neither clearly does not belong

This chapter has a long history, as it went through a considerable number of changes and revisions before taking its present shape. In the course of this history, several people have helped me with it in various ways. Richard Dietz, Max Kölbel and Tim Williamson are definitely among them. But there are many others.

¹ See Horwich 1998, 78–84 and Field 1994, 409–22.

² Williamson 1995, 171.




to it. This reading hides an ambiguity. An object may be a borderline case for a predicate in the sense that it is unclear *simpliciter* whether it belongs to the extension of the predicate. That is, an object may neither paradigmatically have a property nor paradigmatically not have it. For example, if the height of A is 1.75, A is neither paradigmatically tall nor paradigmatically not tall. Thus A may be called a borderline case for 'tall'. Alternatively, an object may be a borderline case for a predicate relative to a certain context in the sense that it is unclear *in that context* whether it belongs to the extension of the predicate. For example, A is a borderline case for 'tall' relative to a certain context if it is unclear in that context whether 'tall' applies to A.

The distinction between borderline situations and borderline objects doesn't really matter as far as borderline objects in the second sense are concerned. Since the characterization of such an object is parasitic on the notion of a borderline situation, it turns out to be trivial that borderline situations arise just in case borderline objects are involved. However, it is important to avoid confusion between borderline situations and borderline objects in the first sense, because there is no reason to think that they must always go together. It is plausible to say that—given the due qualifications—whenever a borderline situation arises, a borderline object in the first sense is involved. But the converse does not hold. Borderline objects in that sense do not necessarily make borderline situations. Suppose that A is 1.75. This does not entail that whenever 'A is tall' is assertively uttered, problems of evaluations arise. It is easy to imagine cases in which the sentence is assertively uttered, yet no unclarity affects the evaluation of the assertion. If A is teacher of a class of kids, and one of them assertively utters 'A is tall' in order to explain to another why A is able to write on a part of the board they can't reach, the assertion may be taken to be correct. If instead 'A is tall' is assertively uttered by someone who is seriously considering A's chances to join a basketball team, the assertion may be taken to be incorrect.

Here only the first reading of 'borderline case' will be adopted. To avoid confusion, the expression will not be used to refer to objects. In the literature the two readings are often mixed, and the ambiguity of the second is often neglected. This is why sometimes philosophers talk as if whenever (and just because) a borderline object in the first sense is involved, a borderline case should arise. And given that borderline objects are as common as paradigmatic objects, this amounts to talking as if borderline cases were as common as unproblematic cases. But such talk is just a theoretical habit that takes us very far from ordinary linguistic practice. In reality, it almost never happens that a speaker assertively utters a sentence and we don't know how to evaluate the assertion. We normally take for granted some way of understanding the sentence according to which the assertion turns out clearly correct or clearly incorrect. Borderline cases almost never arise. As a matter of fact, borderline objects in the first sense are correctly or incorrectly described this or that way depending on the occasion.

The second clarification concerns the expression 'what is said'. Roughly, what is said by uttering a sentence on a certain occasion depends on how the words occurring in the sentence are understood on that occasion, where understanding a word involves grasping its linguistic meaning and, possibly, specifying its reference in the context of utterance. This leaves room for at least two distinct notions. One is that,



when a sentence is assertively uttered by a speaker, there is something that the speaker has in mind. For example, a speaker may use the sentence ‘He is there’ to convey the information that A is in St. Andrews, while another speaker may use it to convey the information that A is in Mexico City. A natural way to express this difference is to say that the two speakers have different things in mind when they utter the sentence. So there is a reading of ‘what is said’—call it *intentional* reading—according to which what is said by uttering a sentence on a certain occasion is a matter of what understanding of the sentence can rightfully be ascribed to the speaker on that occasion.

The other notion is that, when a sentence is assertively uttered by a speaker, there is something to which truth or falsity can be ascribed. The something in question is naturally understood as a specification of the reference of the words occurring in the sentence such that, according to it, either the sentence describes things as they are or it describes things as they are not. Suppose that A is in Mexico City. Then ‘He is there’ turns out true if ‘he’ refers to A and ‘there’ refers to Mexico City, while it turns out false if ‘he’ refers to A and ‘there’ refers to St. Andrews. So there is a reading of ‘what is said’—call it *truth-conditional* reading—according to which what is said by uttering a sentence on a certain occasion is a matter of what understanding makes the sentence evaluable as true or false on that occasion.

It is a naive temptation to put together the two notions. This amounts to thinking that, when a sentence is assertively uttered by a speaker, there is one understanding of the sentence that can rightfully be ascribed to the speaker and involves a specification of the reference of the words occurring in it that makes it evaluable as true or false. Contrary to this temptation, here it will be assumed that the two readings of ‘what is said’ do not coincide. It may happen that a speaker assertively utters a sentence and has something in mind, yet no specification of the reference of the words in the sentence that makes it evaluable as true or false can rightfully be ascribed to the speaker.³


16.2


The thought entertained in this chapter is that borderline cases are cases in which there is no such thing as what is said in the truth-conditional sense. An example of borderline case may help illustrate. Suppose that the grandmother is in the kitchen and the cat is in the living room in such a position that half of its body lies on the mat and the other half lies on the floor. The grandmother makes a guess and assertively utters the sentence

- (1) The cat is on the mat

Apparently, it is hard to tell whether the grandmother’s assertion is correct or incorrect, because it is hard to tell whether the word ‘on’ that occurs in (1) applies to the cat

³ Iacona (2006) deals with cases in which the two readings seem not to coincide.





and the mat. This is not simply due to the fact that the cat is in the position described. It is easy to imagine situations in which (1) is uttered and the cat is in the very same position, yet no problem of evaluation arises. Imagine that the cat got muddy in the garden, and the mat is brand new and very expensive. The grandmother could be concerned about the mat, and wonder whether at least part of the body of the cat lies on it. Or imagine that it is winter and the floor is cold. The grandmother could be concerned about the health of the cat, and wonder whether its whole body lies on the mat.

On the assumption that there are cases in which the two readings of ‘what is said’ do not coincide, this can be regarded as one of them. The grandmother has something in mind. That is, something can rightfully be ascribed to her as the actual understanding of (1). But that understanding does not amount to *one* specification of the reference of the words occurring in (1) that makes (1) evaluable as true or false. Whether (1) is true or false depends on whether or not the pair formed by the cat and the mat belongs to the extension of ‘on’. But the extension of ‘on’ can be specified both in such a way as to include the pair and in such a way as not to include it. As the two alternative situations imagined show, neither of the two possibilities is ruled out by the linguistic meaning of ‘on’. This is to say that there are ways of understanding ‘on’ according to which (1) is true and ways of understanding ‘on’ according to which (1) is false. The point is that the actual understanding of (1) does not decide between the former and the latter. What is said in the intentional sense does not uniquely determine something said in the truth-conditional sense.

The case of the grandmother may be described as a case of *underspecification*, that is, as a case in which the actual understanding of a sentence is not sufficiently specific for the purpose of ascribing truth or falsity to the sentence. For there are at least two ways of understanding (1) that go beyond its actual understanding, in the sense of not being uniquely determined by it. A case of underspecification is not a normal case, as normally the actual understanding of a sentence is sufficiently specific.

The account may be phrased in a more rigorous way by using familiar terminology. Let a *valuation* be an assignment of semantic properties to the sentences of a language that determines definite truth conditions for them. For example, in the case of (1) a valuation will assign an object to ‘the cat’, another object to ‘the mat’, and a set of pairs to ‘on’. Valuations amount to what supervaluationists call ‘precisifications’, that is, ways of making language precise.⁴


A valuation is *admissible* when it corresponds to a legitimate way of making language precise, that is, when it respects the constraints imposed by the linguistic meaning of its expressions. For example, a valuation that assigns a dog to ‘the cat’ will not be admissible. The actual understanding of a sentence—what is said in the intentional sense—may be described as a set of admissible valuations, namely, the set of all the admissible valuations that are not ruled out by that understanding.⁵

Truth in a valuation is defined in the standard way. In the case of a simple sentence, it depends on whether the predicate applies to the objects denoted by the terms. Thus (1) will be true in a valuation v just in case the pair formed by the objects that v

⁴ As in Fine (1975).

⁵ Admissibility is again as in Fine (1975).





assigns to 'the cat' and 'the mat' belongs to the set that v assigns to 'on'. The truth of a complex sentence depends on that of its constituents, in accordance with the usual compositional rules. For example, if (1) is true in v and 'Snow is white' is true in v , then also 'The cat is on the mat and snow is white' will be true in v .

Let us stipulate that two valuations *overlap* on a sentence s when they are alike as far as the truth or falsity of s is concerned. For example, two valuations overlap on (1) if they assign to 'the cat' and 'the mat' the same objects a and b , and they assign to 'on' two sets that differ only in that one includes a pair $\langle c, d \rangle$ formed by a certain piece of paper and a certain table, while the other doesn't include it. In both cases, the pair $\langle a, b \rangle$ belongs to the extension of 'on'.

An *interpretation* of s is a set of admissible valuations that overlap on s . Informally speaking, an interpretation is a way of understanding a sentence that is sufficiently specific for the purpose of ascribing truth or falsity to the sentence on a given occasion. What is said in the truth-conditional sense is an interpretation.



Truth in an interpretation is defined in terms of truth in a valuation. Let α be an interpretation of s . Then, s is true in α if and only if it is true in every valuation that belongs to α . Similarly, s is false in α if and only if it is false in every valuation that belongs to α .


Let an interpretation of s be *compatible* with the actual understanding of s just in case it is a subset of it. The distinction between normal cases and cases of underspecification is then as follows. A normal case is a case in which there is a unique interpretation that is compatible with the actual understanding, namely, the actual understanding itself. A case of underspecification, instead, is a case in which there are at least two interpretations that are compatible with the actual understanding and yield opposite truth values. To see this, suppose that the actual understanding α of s is not an interpretation. Then there are at least two valuations v_1 and v_2 that belong to α but do not overlap on s . Let β be a set whose only member is v_1 , and let γ be a set whose only member is v_2 . β and γ are interpretations compatible with α that yield opposite truth values for s . Inversely, suppose that there are two interpretations β and γ that are compatible with α and yield opposite truth values for s . Then α is not an interpretation. For if we take a member v_1 of β and a member v_2 of γ , we get that v_1 and v_2 belong to α but do not overlap on s .

The underlying thesis of the account is that truth and falsity apply to sentences *relative to interpretations*. On the truth-conditional reading of 'what is said', this means that the things we say are interpretations of the sentences we utter. The distinction between normal cases and cases of underspecification may thus be described in terms of attribution of things said to speakers. Let the condition of attribution be as follows for an interpretation α and a speaker S:

α can be attributed to S when only α is compatible with the actual understanding manifested by S.

There are cases in which something true is said, that is, cases in which an interpretation can be attributed to the speaker and the sentence uttered is true in that interpretation. Similarly, there are cases in which something false is said, that is, cases in which an interpretation can be attributed to the speaker and the sentence uttered is






false in that interpretation. The case of the grandmother belongs to neither of these two categories, as it is a case in which there is no unique thing said. One way to make sense of this is to allow that more than one thing is said, in that more than one interpretation is compatible with the actual understanding. Another way is to hold that nothing is said, in that no interpretation can be attributed to the speaker. The latter seems preferable. First of all, it is quite implausible to hold that a speaker says both something true and something false at the same time. Secondly, although it is sensible to assume that any thing said in the truth-conditional sense is an interpretation, there seems to be no reason in addition to grant that any interpretation is something said in the truth-conditional sense. For example, any set containing one admissible valuation is by definition an interpretation of (1). But we wouldn't describe any such set as something said by the grandmother.


A case of underspecification differs from a normal case in that it involves a conflict of interpretations. The problem that characterizes it concerns the attribution of interpretations, it is not primarily a problem of truth and falsity. There is an obvious sense in which underspecification is compatible with the principle of *bivalence*, according to which truth and falsity are mutually exclusive and jointly exhaustive values. It is the sense in which bivalence holds in any interpretation. By definition, for any interpretation α of a sentence s , either all the valuations that belong to α make s true or they make s false. This entails two claims. One is that either s is true in α or s is false in α : truth and falsity are jointly exhaustive. The other is that it cannot be the case that s is true in α and s is false in α : truth and falsity are mutually exclusive.

16.3

So far it has been suggested that the case of the grandmother is a case of underspecification. Since the case of the grandmother is clearly a borderline case, this entails that at least some borderline cases are cases of underspecification. However, it does not entail that all borderline cases are cases of underspecification. The stronger claim may be justified as follows. Let C be a case in which S assertively utters s , and let α be the actual understanding of s manifested by S . Now suppose that α is an interpretation. Given that bivalence holds relative to interpretations, s is either true or false in α . If s is true in α , then an interpretation in which s is true can be attributed to S . This means that no problem of evaluation arises in C . Similarly, if s is false in α , then an interpretation in which s is false can be attributed to S . This means that no problem of evaluation arises in C . Therefore, if α is an interpretation then no problem of evaluation arises in C . Since C is a borderline case only if some problem of evaluation arises in C , we get that if α is an interpretation then C is not a borderline case. It follows that if C is a borderline case then α is not an interpretation. That is, underspecification obtains in all borderline cases.

We saw that if α is not an interpretation, there are at least two interpretations compatible with α that yield opposite truth values for s . Therefore, the claim that underspecification obtains in all borderline cases entails that if C is a borderline case, there are at least two interpretations compatible with α that yield opposite truth values for






s. A fortiori, it entails that if C is a borderline case, there are admissible valuations that yield opposite truth values for *s*. This amounts to a presupposition that is at the core of semantic theories of vagueness, namely, that borderline cases are cases in which the linguistic meaning of the words occurring in the sentence does not settle the question of whether the assertion is correct or incorrect. For example, a case in which a speaker assertively utters 'A is tall' and we don't know how to evaluate the assertion is a case in which the meaning of 'tall' rules out neither the possibility that the assertion is correct nor the possibility that it is incorrect.

A general tenet of semantic theories of vagueness is that the kind of unclarity that is characteristic of borderline cases depends on the linguistic meaning of the words occurring in the sentences we utter. The argument above shows that this tenet holds, in that it is a necessary condition of borderline cases that the linguistic meaning of the words occurring in the sentence uttered does not determine whether the assertion is correct or incorrect. On the other hand, however, underdetermination in this sense is *not* a sufficient condition of borderline cases. The linguistic meaning of 'on' and the fact that the cat is in such a position that half of its body lies on the mat and the other half lies on the floor do not suffice to make the case borderline. The same goes for the linguistic meaning of 'tall' and the fact that the height of A is 1.75. In order to have a borderline case we need three ingredients: the linguistic meaning of the words occurring in the sentence, the state of affairs, and the actual understanding. The first two do not suffice.

This point may be phrased in terms of the customary distinction between speaker's meaning and semantic meaning, that is, between what the speaker means by uttering certain words, and what those words mean. Since borderline cases are effects of the actual understanding of sentences, they are features of speaker's meaning. Therefore, there is a sense in which vagueness concerns speaker's meaning rather than semantic meaning. Semantic theories of vagueness usually do not recognize this sense. Yet there is also a sense in which vagueness concerns semantic meaning, that in which it amounts to underdetermination of linguistic meaning. The semantic meaning of a vague sentence may be identified with a set of valuations some of which are in conflict. In this sense it is correct to say that borderline cases are grounded on linguistic meaning. Asking which of these two senses comes first, or which is more fundamental, is like asking whether the chicken or the egg comes first. On the one hand, vagueness in speaker's meaning depends on vagueness in semantic meaning, in that the latter is a condition of the possibility of the former. On the other, vagueness in semantic meaning depends on vagueness in speaker's meaning, in that words have linguistic meaning as a result of the way speakers use them.

16.4

The account of borderline cases in terms of underspecification contradicts Williamson's disjunctive moral. In the first place, borderline cases pose no threat to classical logic. For they involve no violation of bivalence. In the second place, the disquotational principle is preserved. Assuming that truth and falsity apply to



sentences relative to interpretations, there are two ways of stating the principle. One is to phrase it as a schema that holds for sentences as they are actually understood and involves ascription of truth relative to interpretations. That is, for any interpretation α , the following biconditional is true relative to α :

(T1) 'p' is true if and only if p

Here 'p' is a substitutional variable, and 'true' is a predicate whose extension varies with α . Borderline cases do not affect the principle, in that the only trouble they give concern the choice of α . The other way is to phrase the principle as a schema that holds for interpreted sentences. For any α , let 'p $^\alpha$ ' stand for 'p' as it is interpreted according to α . The following biconditional is true *simpliciter*:

(T2) 'p $^\alpha$ ' is true if and only if p $^\alpha$


In this case, to see that borderline cases do not affect the principle it must be acknowledged that (T2) follows from a more general schema:

(T3) If 'p' says that p $^\alpha$, then 'p' is true if and only if p $^\alpha$

A borderline case is a case in which 'p' is uttered, but there is no α such that 'p' says that p $^\alpha$. This means that the antecedent of (T3) is not satisfied, hence that (T3) is vacuously true.⁶

In the third place, the account does not entail that vagueness is an epistemic phenomenon. According to epistemicism, the meaning of the words occurring in a sentence determines a truth value for the sentence, in borderline cases just as in normal cases. This is to say that in the case of the grandmother the extension of 'on' is such as to make (1) true or false. The idea is that words have sharp boundaries, but we don't know exactly where these boundaries lie. That is, given any cat and any mat, either the pair formed by them belongs to the extension of 'on' or it doesn't. The fact is that we don't know exactly the borders of that extension, hence it may happen that we don't know, of a certain cat and a certain mat, whether or not the pair formed by them belongs to it. The account in terms of underspecification entails nothing like that. Truth and falsity apply to (1) relative to interpretations, so only relative to this or that interpretation the extension of 'on' is such as to make (1) true or false. Since borderline cases are cases in which there is no such thing as 'the' correct interpretation, there is no such thing as the interpretation that makes (1) true or false. So there is nothing to be ignorant of.

⁶ (T3) provides an argument against the hypothesis—whose plausibility is called into question in §2—that borderline cases are cases in which more than one thing is said. As it is shown in Andjelkovic and Williamson 2000, 225–6, (T3) entails a principle of *uniformity* to the effect that everything said by a sentence on a certain occasion has the same truth value: if 'p' says that p $^\alpha$ and 'p' says that p $^\beta$, then p $^\alpha$ if and only if p $^\beta$. For suppose that there are two interpretations α and β such that 'p' says both that p $^\alpha$ and that p $^\beta$. By instances of (T3), this yields that 'p' is true if and only if p $^\alpha$ and that 'p' is true if and only if p $^\beta$. It follows that p $^\alpha$ if and only if p $^\beta$. The principle of uniformity rules out the hypothesis that there are cases in which 'p' is uttered, and there are two interpretations α and β such that 'p' says both that p $^\alpha$ and that p $^\beta$ but α and β yield opposite truth values.



This provides a response to an argument that is a mainstay in Williamson's defence of epistemicism. The argument goes as follows. Suppose that 'A is tall' is uttered and we are in a borderline case. If one accepts a plausible unknowability principle for borderline cases, one gets that we don't know whether A is tall. But if one assumes classical logic and the disquotational principle, one also gets that 'A is tall' is either true or false. This means that there is something to be known: 'the speaker's ignorance has an object'.⁷ The response is that it must not be taken for granted that classical logic and the disquotational principle entail that 'A is tall' is either true or false *simpliciter*. On the assumption that truth and falsity apply to sentences relative to interpretations, bivalence can be shown to hold relative to interpretations, in accordance with classical logic and the disquotational principle. Therefore, what one gets is that, *in any interpretation*, 'A is tall' is either true or false. This is consistent with the hypothesis that the speaker's ignorance has no object in the case considered. For in that case there is no such thing as the correct interpretation of 'A is tall'.⁸

16.5

Showing that there is a coherent view that preserves classical logic and the disquotational principle without being epistemicist wouldn't be enough if the view were implausible. For mere coherence does not justify acceptance. This is what Williamson seems to think of a view that looks pretty much like that outlined in the previous sections:

At least one view does combine classical logic and (T+) with the denial that vagueness is an epistemic phenomenon in a way that seems at any rate not formally inconsistent. This is the view that vague sentences do not say that anything is the case, in borderline cases if not elsewhere. Thus, if 'b' is a borderline case for 'bald', to say 'b is bald' or 'b is not bald' is to make an utterance without propositional content. One says nothing by uttering either sentence, so neither is true. Thus, the corresponding antecedents of (T+), in 'in c "b is bald" says that P' and 'in c "b is not bald" says that P', fail whatever legitimate substitution is made for 'P', and (T+) holds vacuously. For the same reason, (T) would have no relevant instance. Thus, there would be no relevant true sentence, and therefore nothing for speakers of the language to be ignorant of. [...] No attempt will be made here to argue against that extreme view. Many of the philosophers who wish to accept orthodoxy while denying that vagueness is an epistemic phenomenon are willing to instantiate (T) with vague sentences even in borderline cases. They allow that 'b is bald' is true if and only if b is bald; 'b is bald' says that b is bald, even if b is borderline for 'bald'.⁹

Here (T+) is a principle equivalent to (T3), (T) is the disquotation schema, and 'orthodoxy' stands for the combination of classical logic with (T). The account outlined in the previous sections seems to fit the description. There is a sense in which

⁷ Williamson 1995, 174–5.

⁸ Obviously, here the assumption is that Williamson's use of 'borderline case' does not differ from ours.

⁹ Williamson 1995, 173.



one says nothing by uttering a sentence in a borderline case, namely, the truth-conditional sense. In this sense it is correct to say that a principle equivalent to (T+) holds vacuously, in that the antecedent of (T3) is not satisfied.


Why should the view be ‘extreme’? The alleged motivation is that an utterance of ‘b is bald’ in a borderline case would have no content. Consequently, the disquotation schema wouldn’t be instantiated, and we couldn’t say that ‘b is bald’ is true if and only if b is bald. But that isn’t so. It is plausible that an utterance of ‘b is bald’ in a borderline case has content. The obvious sense in which it has content is that some understanding of the sentence can rightfully be ascribed to the speaker. In other words, such a case differs from one in which, say, ‘He is there’ is written on the board to make a point of grammar. However, the view does not entail that ‘b is bald’ lacks content in *that* sense. For that sense is the intentional sense, and the view does not deny that something is said in the intentional sense.¹⁰ Similarly, it is plausible that (T) is instantiated with sentences uttered in borderline cases, as it is quite natural to say that ‘b is bald’ is true if and only if b is bald. But again, the view does not deny it. The disquotational principle can be phrased in terms of a schema—(T1)—that holds for sentences as they are actually understood, and these obviously include sentences uttered in borderline cases. So the view is not extreme, and perhaps some attempt should be made to argue against it.

To appreciate its plausibility, it is important to distinguish it from another view that is indeed extreme, namely, that according to which only precise expressions have meaning, so whenever we utter a sentence containing vague expressions we say nothing at all.¹¹ What the account of borderline cases outlined requires is simply that there are cases in which we say nothing *sufficiently precise*. This is compatible with there being cases—the normal ones—in which we say sufficiently precise things. For example, even if there are cases in which ‘on’ is understood in such a way as to say nothing sufficiently precise by uttering (1), most of the time ‘on’ is understood in ways that are sufficiently precise for the descriptive goals that guide our use of the sentences containing it.

This is not to say that most of the time ‘on’ is understood in a completely precise way. The use of a predicate almost never involves complete specification of its extension. Rather, it involves a *partial* specification of it, in that the only part of extension that matters on each occasion concerns the objects that are salient on that occasion. In other words, the use of a predicate on a given occasion determines a set of valuations, the set of all the valuations that share the part specified. Thus if one uses (1) to describe a cat *a* and a mat *b*, one takes the extension of ‘on’ to include the pair $\langle a, b \rangle$. But there may be another pair $\langle c, d \rangle$, such that it is left indeterminate by that use whether $\langle c, d \rangle$ belongs to the extension of ‘on’. Therefore, that use determines a set that includes at least two valuations: in one of them the extension

¹⁰ García-Carpintero 2007, draws a distinction that is analogous at least as far as Williamson’s appeal to intuitions is concerned.

¹¹ This is the view called ‘Nihilism’ in Williamson 1994, ch. 6. The same view is considered in McGee and McLaughlin 1995, 215.



of 'or' includes $\langle a, b \rangle$ but not $\langle c, d \rangle$, in the other it includes both $\langle a, b \rangle$ and $\langle c, d \rangle$. If we call completely specified a way of understanding a sentence that provides for all possible cases, then sufficient specification does not mean complete specification. A completely specified way of understanding a sentence amounts to a single valuation rather than to a set of valuations.

16.6


The foregoing sections show how borderline cases can be accounted for in accordance with the thesis that truth and falsity apply to sentences relative to interpretations. This section shows that the thesis provides a straightforward solution to the paradox of the sorites. In its original form, the sorites says that if the removal of one grain from a heap always leaves a heap, the successive removal of every grain still leaves a heap. The argument goes as follows:


- (2) 1000 grains make a heap
- (3) For every n , if n grains make a heap then $n - 1$ grains make a heap
- (4) 0 grains make a heap

The inference from (2) and (3) to (4) seems legitimate, (2) and (3) seem acceptable, but (4) seems unacceptable. Therefore, it is natural to think that there is something wrong with the argument. To provide a solution to the paradox is to say what exactly is wrong.

The thesis that truth and falsity apply to sentences relative to interpretations points to the following definitions of validity and soundness. An argument is *valid* if and only if, necessarily, every interpretation that makes its premises true also makes its conclusion true. The criterion involved is the classical one, that according to which validity is necessary truth preservation. But since the premises and conclusion of an argument can be true only relative to interpretations, truth can be necessarily preserved only relative to interpretations. Similarly, an argument is *sound* in an interpretation if and only if it is valid and all its premises are true in that interpretation. Again, the criterion involved is the classical one, that according to which a sound argument is a valid argument whose premises are true. But since the premises of an argument can be true only relative to interpretations, the argument can be sound only relative to interpretations.

These two definitions tell us what is wrong with the sorites. Although the argument is valid, there is no interpretation that makes it sound. For there is no interpretation in which (3) is true. An interpretation of (3) is a way of understanding (3) relative to which truth or falsity can be ascribed to it. Since (3) is equivalent to a list of conditionals, the interpretation must allow ascription of truth or falsity to each of the conditionals in the list, hence it must allow ascription of truth or falsity to the antecedent and to the consequent of each of the conditionals in the list. This entails that, for each of the collections of grains featuring in the series that goes from 1000 to 0 grains,






it has to be specified whether or not it is a heap. Therefore, an interpretation of (3) involves a way of understanding ‘heap’ according to which there is a cut-off point in the series, that is, there is a number n such that a collection of n grains belongs to the extension of ‘heap’ while a collection of $n - 1$ grains does not belong to it.


The sorites draws its appeal from the fact that we normally use ‘heap’ *without* specifying its extension as required by (3). Thus, (2) seems acceptable in that ‘heap’ is normally understood in such a way that a collection of 1000 grains belongs to its extension. Similarly, (4) seems unacceptable in that ‘heap’ is normally understood in such a way that a collection of 0 grains does not belong to its extension. But such ways of understanding ‘heap’ do not involve complete specification of its extension. In particular, they do not involve a delimitation of that extension sensitive to differences of one grain. Normally, when we understand ‘heap’ in such a way that a collection of n grains belongs to its extension, we do not have in mind a specification which prescribes that a collection of $n - 1$ grains does not belong to it. Therefore, we are apt to exclude that a collection of n grains is a heap *but* a collection of $n - 1$ grains is not a heap. This is why (3) seems acceptable. In other words, what makes the existence of a cut-off point for ‘heap’ unwelcome is that we normally do not specify such a point. Whenever we use the word to describe a certain object, we take for granted that no cut-off point lurks in the vicinity of that object, namely, that relevantly similar objects may equally be described in the same way.¹²

16.7

The account of vagueness outlined in this chapter—call it the *underspecification view*—may be contrasted with two similar accounts that are well known. The first is standard supervaluationism. The underspecification view substantively differs from standard supervaluationism, as the latter does not contradict Williamson’s moral. The basic idea of supervaluationism is that the vagueness of natural language consists in its capacity in principle to be made precise in more than one way. Following this idea, the method of supervaluations is adopted in order to deal with sets of precisifications of the language. A supervaluation is an assignment of truth values based on a quantification over assignments of truth values relative to precisifications. A sentence is ‘supertrue’ if it is true on all precisifications, ‘superfalse’ if it is false on all precisifications, and neither otherwise. Supervaluationism identifies truth with supertruth. This is why it obeys Williamson’s moral. In the first place, supertruth does not conform to the disquotational principle. If a sentence is true on some precisifications and false on others, the biconditional obtained by plugging the sentence in the disquotation schema is not supertrue. In the second place, supertruth violates bivalence, hence

¹² Assuming that the use of ‘heap’ on a given occasion is guided by what is psychologically or conversationally salient on that occasion, one may say that in normal circumstances no cut-off point is located within the area of salience. This, however, is not quite the same thing as to say that a cut-off point exists but is located somewhere outside that area, as suggested in Fara 2000.





classical logic. Moreover, as it has been emphasized by its critics, supervaluationism fails to preserve some classical principles about consequence and rules of inference. This is due to the definition of validity as necessary preservation of supertruth, that is, the definition according to which an argument is valid just in case necessarily, if its premises are supertrue then its conclusion is also supertrue.¹³

The underspecification view shares with supervaluationism its basic idea. It is plausible to say that a vague language is a language that in principle can be made precise in more than one way. This is couched by the assumption that a sentence allows different admissible valuations. However, the underspecification view differs from supervaluationism in that it does *not* identify truth with supertruth. Truth in an interpretation conforms to the disquotational principle and does not violate bivalence.¹⁴ Moreover, the underspecification view has no problem with consequence and rules of inference, in that it does not involve a supervaluational definition of validity. The latter says that an argument is valid just in case, necessarily, if its premises are true on all precisifications then its conclusion is true on all precisifications, while the definition given in §6 says that an argument is valid just in case, necessarily, in all interpretations, if its premises are true then its conclusion is true. The difference between the two definitions may be seen as a difference in the scope of the quantification over precisifications or interpretations.¹⁵

The second account to be considered is the non-standard version of supervaluationism advocated by Van McGee and Brian McLaughlin, and then adopted by Cian Dorr. The version at issue is non-standard in that it does not identify truth with supertruth. McGee and McLaughlin claim that a distinction must be drawn between ‘truth’ and ‘definite truth’, where the former is defined in accordance with the disquotational principle and classical logic, while the latter is characterized by using a supervaluational model-theoretic apparatus. Given that this characterization of definite truth is not epistemic, their view is like the underspecification view—and unlike standard supervaluationism—in that it contradicts Williamson’s moral.¹⁶

But some significant differences remain. In the first place, it is not clear whether the actual understanding of a sentence plays some role in the determination of definite truth. McGee and McLaughlin say at a certain point that ‘the thoughts and practices of the speakers of the language, together with the non-linguistic facts, pick out a set of sentences as definitely true’.¹⁷ However, it is not clear whether this entails that what a speaker has in mind on a certain occasion may contribute to determine what


¹³ See Williamson 1994, 146–53. The problem concerns supervaluationism in its standard version. But other definitions may be adopted. See Keefe 2000 and Varzi 2007.

¹⁴ The simple fact that truth in an interpretation is defined in terms of a quantification over valuations does not mean that it is a form of supertruth, at least not in the sense that matters here. For the definition does not allow for a third status between truth and falsity.

¹⁵ Williamson 1994, 147–8, Keefe 2000, §3, and Varzi 2007 spell out this distinction. The two kinds of definitions are called ‘global’ and ‘local’ respectively.

¹⁶ McGee and McLaughlin 1995, Dorr 2003. Williamson 1997, however, questions that the characterization of definite truth in McGee and McLaughlin 1995 is not epistemic.

¹⁷ McGee and McLaughlin 1995, 227.




is definitely true on that occasion. If it doesn't, then in the case of the grandmother the view must be that the unclarity that affects the evaluation of (1) is independent of what the grandmother has in mind, although it may depend on the thoughts and behaviour of the linguistic community as a whole. This is to say that the unclarity at issue is due entirely to the fact that the linguistic meaning of 'on' and the position of the cat are as they are, contrary to what is assumed here.¹⁸

In the second place, even supposing that McGee and McLaughlin do take into account the actual understanding of a sentence, or that their view may be adjusted in such a way as to take it into account, a difference remains. When McGee and McLaughlin talk of truth and definite truth, they seem to have in mind what is said in the intentional sense, whereas here the bearer of truth and falsity is what is said in the truth-conditional sense. This is in part a matter of focus. The underspecification view could as well be phrased in terms of the intentional reading of 'what is said'. Borderline cases could be described as cases in which a speaker asserts a unique thing that is not evaluable. This would not essentially differ from saying that borderline cases are cases in which the thing asserted is neither definitely true nor definitely false, given that 'evaluable' would mean 'evaluable as true or false'. However, there is one crucial respect in which the difference is not simply a matter of focus. McGee and McLaughlin ascribe truth *simpliciter* to what is said in the intentional sense. This entails that the things we say in borderline cases are true or false *simpliciter*. By contrast, even if the underspecification view were phrased in terms of the intentional reading of 'what is said', such ascription would be ruled out. For its main thesis would be that the things we say are evaluable as true or false only relative to interpretations.

References

- Andjelkovic M. and Williamson, T. (2000), 'Truth, falsity, and borderline cases', *Philosophical Topics* 28, 211–44.
- Dorr, C. (2003), 'Vagueness without ignorance', in Hawthorne, J. and Zimmerman, D., eds., *Philosophical Perspectives* 17, Blackwell, 83–114.
- Fara, D. G. (2000), 'Shifting sands: an interest-relative theory of vagueness', *Philosophical Topics* 28, 45–81. (Originally published under the name "Delia Graff").
- Field, H. (1994), 'Disquotational truth and factually defective discourse', *Philosophical Review* 103, 405–52.
- Fine, K. (1975), 'Vagueness, truth and logic', *Synthese* 30, 265–300.
- García-Carpintero, M. (2007), 'Bivalence and what is said', *Dialectica* 61, 167–90.
- Horwich, P. (1998), *Truth*, Oxford University Press, Oxford.
- Iacona, A. (2006), 'True in a sense', *Grazer Philosophische Studien* 72, 141–54.
- Keefe, R. (2000), 'Supervaluationism and validity', *Philosophical Topics* 28, 93–106.
- McGee, V. and McLaughlin, B. (1995), 'Distinctions without a difference', *The Southern Journal of Philosophy* 33 (Supplement), 203–51.
- Varzi, A. (2007), 'Supervaluationism and its logics', *Mind* 116, 633–75.

¹⁸ This is what they seem to think when they say things such as 'Harry is bald' is definitely true if Harry is 'definitely bald', 210. Similar considerations hold for Dorr 2003.



Williamson, T. (1994), *Vagueness*, Routledge, London.

— (1995), 'Definiteness and knowability', *The Southern Journal of Philosophy* 33 (Supplement), 171–91.

— (1997), 'Imagination, stipulation and vagueness' in Villanueva, E. (ed.), *Philosophical Issues* 8, Ridgeview, Atascadero (CA), 214–28.

