

Symposium

on

Understanding Truth

Reply

to

Anil Gupta, Matthew McGrath, Jamie Tappenden and Timothy Williamson

by

Scott Soames

Philosophy and Phenomenological Research
Volume LXV, No. 2, 2002

**Reply to Anil Gupta, Matthew McGrath, Jamie Tappenden
and
Timothy Williamson**

Reply to Gupta

Anil Gupta gives a good explication of what I mean by a partially defined predicate, as opposed to a completely defined partial-predicate. He then asks whether we should accept the thesis that *true* is partially defined, like my example *smidget*. He suggests that we shouldn't. However, his negative thesis is ambiguous.

Interpretation 1

Unlike *smidget*, *true* is not a partially defined predicate.

Interpretation 2:

True is a partially defined predicate, but the rules that govern it are importantly different from those governing *smidget*.

In the end it becomes clear that Gupta's remarks support at most the second interpretation. Even then it remains to be seen just how different, and just how successful, his alternative is.

His first point is that *true* exhibits pathologies that *smidget* doesn't. If *smidget*, is undefined for Charlie, then the sentence *Charlie is a smidget* is undefined, and there is no basis for accepting either it or its negation. There is no pathology here; it is simply a case in which a sentence and its negation must both be rejected. With *smidget* there is no paradoxicality analogous to Liar sentences and no circularity corresponding to Truth Tellers. Gupta concludes that *true* and *smidget* must be very different.

He is right; the two predicates are different. But this is no criticism of my characterization. The pathologies found with *true* result from the fact that it applies to sentences containing it, whereas *smidget* does not. On the version of the Kripke model I

suggest, the extension and antiextension of *true* are built up cyclically. First the truth rules are applied to sentences that don't contain it. After adding these truths and untruths to the extension and antiextension, one applies the rules again – this time to sentences containing the truth predicate. Further applications continue indefinitely, with the result that more and more sentences are characterized as true, or untrue. However, Liars and Truth Tellers never get evaluated, because their evaluation involves a kind of circularity. They can be characterized as true, or untrue, at any stage of the process only, if they have already been characterized at an earlier stage. Since they are not characterized at the first stage, they never are. Consequently, the truth predicate is partially defined. Unlike *smidget*, which is governed by rules designed to be explicitly and transparently partial, *true* is governed by rules that we discover to be partial because, for certain sentences, any attempt to apply them will be circular, and will yield no result.

As for the different pathologies of Truth Tellers and Liars, that too can be explained. As Kripke points out, one can consistently supplement the truth rules with a stipulation characterizing Truth Tellers as true, or a stipulation characterizing them as untrue. Once this is done, the sentences will cycle through the normal process for assigning truth-values, and one's initial decision will be confirmed. With Liar sentences this is not so. Any decision to treat them as true, or untrue, will lead, via the original rules, to contradiction. As a result, Liars differ from Truth Tellers in that it is **not** possible to consistently supplement the original rules governing the predicate with a stipulation assigning them a truth-value. Hence, my version of Kripke can accommodate the different pathologies of Truth Tellers and Liars.

Gupta's second criticism is that my model of a language containing its own truth predicate cannot be applied if the language contains a certain biconditional operator. In my book I introduced the *just in case* operator, symbolized with three horizontal lines. A sentence $A \equiv B$ is true when both A and B are true, when both are false (not true), and when both are undefined.¹ Otherwise $A \equiv B$ is false (not true). This is the use of *if and only if* that Gupta talks about. As I pointed out in the book, if a language with its own partially defined truth predicate contains this connective plus the means of freely referring to its own sentences, then one can define a predicate in the language that applies to all and only determinate truths of the language. Since this is a totally defined predicate, we can use it to generate a contradiction that can't be blocked by appeal to partiality. The lesson is that no language, with certain minimal expressive capacities, can contain its own partially defined truth predicate plus the *just in case* connective.

Gupta speaks as if this were a regrettable limitation. He says, "*We grant that there is no possibility of a universal language, but in asking for an account of truth in a language with 'iff', we are not asking for a universal language. We can grant further that certain kinds of Tarskian hierarchies are unavoidable, but this does not mean that the invocation of a Tarskian hierarchy is acceptable for 'iff'.*" I gather from these remarks that Gupta thinks that a language with the means of describing, quantifying over, and referring to its own sentences in unrestricted ways **can** contain both its own truth predicate and the *just in case* connective.

But that can't be right. As Gupta points out, *Snow is black $\equiv c$ is true* is a Liar sentence, when c refers to that sentence itself. But now, the only assumptions needed to generate a contradiction are; (i) that a sentence is true (not true) \equiv it is in (not in) the

extension of *true*, (ii) that an atomic sentence *a is P* is true (not true) \equiv a refers to something in (not in) the extension of P, (iii) that *Snow is black* is not true, (iv) that a sentence *S \equiv Q* is true iff both S and Q are true, both not true, or both are undefined (otherwise it is not true); and (v) that every sentence is either true, not true, or undefined. These assumptions are undeniable, no matter whether the truth predicate is partial or not. The only way to avoid contradiction is to deny that a language could satisfy these assumptions by containing its own truth predicate plus the *just in case* connective. But if such a language is impossible, then it is no criticism of my model that it doesn't do the impossible. Nobody else's does either, Gupta's included.ⁱⁱ

Gupta's final criticism can be paraphrased thus: *true* is partially defined, but the rules governing it are not like the rules governing *smidget*; with *true* the rules **do** give necessary and sufficient conditions for the truth or untruth of each sentence, but because they are circular, we end up with a partially defined predicate after all. Thus, Gupta says, "*the rules governing 'true' are **not** silent on the classification of any object, and they **do** supply necessary and sufficient conditions for an object to fall under 'true'*". Nevertheless, he adds, some sentences – e.g. Liars – "*should neither be deemed true nor deemed untrue*". Hence, he ends up treating the truth predicate as partial after all.

These remarks should seem perplexing. If, as Gupta says, the rules for *true* are **not** silent about the classification of any sentence, if they provide necessary and sufficient conditions for the truth or untruth of every sentence, then they will assign Liar sentences truth-values. But if Liar sentences are assigned truth-values, how is it that "*they should neither be deemed true nor deemed untrue*"?

To answer this question, we must go beyond Gupta's explicit remarks. Here is my interpretation: On Gupta's view, we get into trouble if we think that there is just one set of rules governing the truth predicate. Really there are two. The first set assigns prima-facie truth-values. These are familiar Tarski-style rules. They apply to all sentences, and the prima-facie truth-values they assign to most sentences later become their real truth-values. However for Liars and Truth Tellers this is not so. The Tarski-style rules characterize Liars as both prima-facie true and prima-facie untrue. Truth Tellers are characterized as prima-facie true, or as prima-facie untrue, depending on which arbitrary assumption one starts with.

One can think of the Tarski-style rules as being completely reliable in their assignment of truth-values to ordinary sentences, as being unreliable in their assignments to Liars, and as being of questionable reliability in their treatment of Truth Tellers. Since the assignments they make are sometimes arbitrary, and sometimes conflicting, they must be understood as merely tentative. Because of this, a further set of instructions is needed to resolve conflicts among the prima-facie truth-values. Details aside, the rules assigning **real** truth-values tell us that a sentence is true, if it is **stably** characterized as prima-facie true, and it is not true if it is stably characterized as prima-facie untrue. Since Liars are unstably characterized, first as true, then as not true, and so on, the rules assigning **real** truth-values are partial, just like those for *smidget*. Unlike the prima-facie truth rules, the real-truth rules are transparently and obviously silent about certain sentences.

If I am right, then Gupta's picture of truth as partially defined has important similarities with my version of Kripke's picture. There are, of course, significant formal and philosophical differences as well. One important part of Gupta's thinking, as I see it,

is that the ordinary rules we learn when a predicate is introduced or mastered, may sometimes turn out to be defective, in the sense of leading to inconsistencies when applied with full generality. When this is so, it is not always necessary to throw out the predicate and start over from scratch. Instead there may be a process of adjustment, the effect of which is to trade inconsistency for partiality. Gupta seems to think that we should understand the partiality of the truth predicate as arising in this way. In addition, he seems to believe that the process of adjustment needed in the case of the truth predicate might be a general one that applies to other inconsistent definitions as well. If so, then he may be onto a fruitful line of investigation. How successful it will ultimately prove to be in shedding light on the Liar remains an open question. However, if I am right in characterizing the philosophical import of his formal work, then his approach to the partiality of the truth predicate has much in common with my version of Kripke's. The dispute between them is a family quarrel.

Reply to McGrath

I next turn to the comments of Matthew McGrath, the first of which concerns the attitude of rejection. In my view there are cases in which the correct attitude to take toward a proposition and its negation is to reject both. McGrath worries about what this comes to.

“If I reject ‘p’, what have I done? I have done more than merely not accept ‘p’. Presumably, if I am sincere, I have expressed my disbelief that p. But what grounds are there for distinguishing a disbelief that p from a belief that not-p? If there were not distinction, then if I were to reject both ‘p’ and ‘not-p’ sincerely, I would be expressing beliefs in contradictory propositions.”

I will here accept McGrath's terminology and stipulate that to disbelieve a proposition p is to entertain p , to judge that one is not warranted in accepting it, and to decline to accept it on that basis. Standardly, when we believe the negation of p , we disbelieve p in this sense. But the converse need not hold. Even when there is no question of p 's being undefined, we may think that there is insufficient available evidence to warrant accepting either p or its negation, and so disbelieve both. The case in which p and its negation are both undefined is a subcase of this situation. In this special case, even a complete grasp of all relevant facts would not provide sufficient evidence to warrant accepting either proposition. Hence, the correct response is to disbelieve both.

McGrath questions this. He suggests that disbelief, like belief, has a "mind-to-world" direction of fit, and that because of this disbelieving both a proposition and its negation is just as bad as believing them. He says, "*Isn't a disbelief that p correct if not- p , just as a belief that p is correct if p ? If so, then in disbelieving a proposition and its negation, I guarantee that one of the two disbeliefs is incorrect.*" But this conclusion doesn't follow. What McGrath really needs is not the relatively innocuous principle *if not- p , then to disbelieve that p is correct*, but the converse, *if to disbelieve p is correct, then not- p .*ⁱⁱⁱ However, this principle is clearly unacceptable.^{iv}

McGrath's next comment concerns logical inference and logical consequence. He says that I must reject the claim that $P \supset Q$ is a logical consequence of $P \rightarrow Q$, where ' \rightarrow ' is the conditional counterpart of my "just in case" biconditional connective, ' \equiv ', and ' \supset ' is the strong Kleene version of the material conditional. This is correct, if (1) and (2) govern the two conditionals, and (3) characterizes logical consequence in a system that allows undefined sentences.^v

		Q		
		T	F	*
1.	$P \rightarrow Q$	T	F	*
	P	T	F	F
		F	T	T
		*	F	T

		Q		
		T	F	*
2.	$P \supset Q$	T	F	*
	P	T	F	*
		F	T	T
		*	*	*

- (3) B is a logical consequence of A iff every model in which A is determinately true (i.e. defined and true) is a model in which B is determinately true.

Under these assumptions, $P \supset Q$ is a not logical consequence of $P \rightarrow Q$ -- since, when both P and Q are undefined, $P \rightarrow Q$ is determinately true, whereas $P \supset Q$ is undefined.

Should this observation cause concern? McGrath points out that it is a consequence of it that Conditional Proof is not a valid rule of inference. Since if it were, we could derive $P \supset Q$ from $P \rightarrow Q$ by first assuming P, then deriving Q by the valid rule of Modes Ponens, and then using Conditional Proof to get the material conditional from $P \rightarrow Q$ alone. A similar point holds for Reductio Ad Absurdum. Hence neither rule is valid in the sense that conclusions derived using it must be determinately true, if the premises are (though other classical rules are valid in this sense). This ought to have been evident anyway, since when P is undefined, so is $P \supset P$ even though P is a logical

consequence of itself, and $P \vee \sim P$ is undefined even though a contradiction can be derived from its negation.

McGrath says that because of this I have to give up important parts of classical logic. That sounds more serious than it really is. Since Conditional Proof and Reductio Ad Absurdum lead to trouble when undefined sentences are involved, restricting their application to fully-defined cases strikes me as no real loss. However the conceptual issues here are worth exploring, and before trying to decide whether giving up parts of classical logic because of undefined sentences is bad, we should first determine what giving up classical logic on this basis would amount to.

Let us start with the classical validity, as applied to sentences. Where S is **any** classically valid sentence schema (involving the usual logical operators) – even $P \supset P$ -- we must give up the claim that all instances of S are true, when dealing with undefined sentences. Nevertheless, (4a) and (4b) remain intact.

- 4a. All instances of classically valid schemata are either true or undefined.
- b. When all subformulas are defined (have determinate truth-values), instances of the classically valid schemata are all determinately true.

With this in mind, we can characterize a sentence as logically valid iff it is true in all models in which it gets assigned a determinate truth-value, and we can characterize a schema as logically valid iff all instances of it with determinate truths value are true.

With this revised understanding all classically valid sentences and schemata remain valid.

Next consider logical consequence. (3) specifies one natural characterization, which I take McGrath to have been assuming. However, there is a slightly different definition which is also quite natural:

5. B is a logical consequence of A iff every model in which A and B are assigned determinate truth-values is one in which B is determinately true, if A is.

This is like the standard definition in that when we say that B is a logical consequence of A because it is true in every model in which A is true, we restrict ourselves to models that interpret both A and B. In particular if B includes vocabulary that A does not, we exclude from consideration models of A that don't interpret that vocabulary. (This is assumed even in (3).) Suppose we extend this idea by thinking of a sentence that is undefined as lacking something that can be provided by an interpretation. Then in assessing whether B is a logical consequence of A, it is natural to adopt (5) and restrict ourselves to models in which both get assigned truth-values. On this understanding, the extension of the logical consequence relation remains the same as it is in classical logic.

The upshot of this is that there is a natural sense in which allowing for undefined sentences doesn't change the results of classical logic at all (though it does change their interpretation). There is, of course, another natural sense, characterized by acceptance of (3), in which there are some changes in the extension of the logical consequence relation, and in the rules of one's proof theory. Nevertheless, it is not obvious that we must choose one of these alternatives over the other. Both conceptions of the relevant logical notions are legitimate when dealing with undefined sentences and both may be developed, depending on one's purposes. Hence, I don't see a cause for concern.

McGrath's third point involves the connection between partially defined predicates and the idea that undefined sentences must be rejected. Regarding the smidget example, he says that it seems to him wrong to **reject** the claim that an adult of intermediate height is a smidget. Better, he thinks, simply to **withhold** judgement on it.

*“But focusing on an object o named by α for which we know ‘smidget’ is undefined, it seems more likely that we would rather withhold judgement on α is a **smidget** and its negation. Our definition, after all, is partial; it says nothing about such cases. We would be going beyond that definition to reject the claims in question.”*

I disagree. The difference between withholding and rejecting comes to nothing, if one adds to withholding the recognition that because the predicate *smidget* is undefined for o , no further empirical information could possibly decide the question. In cases like this, it is correct to reject a proposition and its negation because it is correct to decisively withhold judgement on them, in the sense of refraining from accepting either one on the basis or current evidence, while recognizing that no appeal to additional empirical facts could warrant revising one’s judgement.^{vi}

McGrath next imagines that we later **change** the definition of smidget. He says:

“Imagine we later decided to include some previously undefined cases in the extension of ‘smidget’. Would we have to regard ourselves as then making claims we previously rejected, or rather as making claims about which we were previously silent?”

The answer is that by changing the meaning (semantic content) of *smidget* we would be changing the propositions expressed by sentences containing it, with the result that in accepting or assertively uttering those sentences we would be endorsing or making claims different from any we had previously considered. These claims would not be ones we had earlier rejected, or even entertained.

McGrath's fourth point concerns the property expressed by *smidget*. According to me the property, like the predicate, is partially defined. McGrath has misgivings.

“If ‘smidget’ latches on to a property, it latches on to one that existed prior to the introduction of the word. Soames denies it latches onto a “totally defined” property, that is, a property the either holds or fails to hold of every object. So he concludes, in effect, that it must latch onto a “partially defined” property. But how does it do that? The partial definition of ‘smidget’ does not fix uniquely on a totally defined property, since there are many such properties that satisfy the conditions given. Yet it isn’t clear how it could fix uniquely on a partially defined property, either. If ‘smidget’ fixes on a property, it does so in virtue of the property’s satisfying the conditions stipulated above. But these conditions don’t fix on any unique property, whether totally or partially defined.”

The main complaint here is that the conditions used to introduce the predicate *smidget* don't uniquely determine any property, since those conditions could be satisfied by many different choices of the property expressed by the predicate. There is certainly a sense in which that is right. But the same is true of total definitions. Suppose, for example, that I introduced the predicate *totalsmidget* with the following stipulation: *For all x, x is a totalsmidget iff x is an adult human being less than 4 feet tall.* Following McGrath, we observe that the conditions expressed by the sentence used in giving this definition will be satisfied by any choice of a property to be expressed by *totalsmidget* that is materially coextensive with the phrase on the right-hand side. Clearly, there are many such properties. Nor will we get a unique property even if we modalize the definition.

Properties are like propositions in that necessarily equivalent ones may nevertheless be distinct. Thus, even if the definition were modalized there would still be many choices of the property expressed by *totalsmidget* that would satisfy the conditions expressed by the sentence used in giving the definition.

So do we **never** succeed in identifying a unique property when we give such a definition? Surely not. When we use conditionals or biconditionals to give definitions of the kind we have been considering, we do so with the intention of getting across something stronger than that which is literally expressed by the sentences used. Usually, when we give such a definition it is with the understanding that there is nothing more to the meaning of the relevant term than that which we have indicated. In the case of a total definition – *For all x, x is P iff x is D* --there is nothing more to determining that an object o has (doesn't have) the property expressed by P than determining that o has (doesn't have) the property expressed by D. (Similarly, for establishing, demonstrating, or supporting the claim that o has, or doesn't have, the property expressed by P.) In the case of a partial definition, we have separate clauses for positive and negative cases – *For all x, if x is D_P, then x is P* and *For all x, if x is D_N, then x is not P*. As a result, the corresponding presumption is that there is nothing more to determining that an object o has the property expressed by P than determining that o has D_P, and there is nothing more to determining that o doesn't have the property expressed by P than determining that o has the property expressed by D_N. (Similarly, for establishing, demonstrating, or supporting the claim that o has, or doesn't have, the property expressed by P.) With such presumptions in place we can use partial definitions of the sort governing *smidget* to uniquely pick out partial properties.

Finally, McGrath raises a point about whether undefinedness exists only “in language” or also “in the world.” Certainly it exists in language, since some sentences are undefined. However, McGrath suggests that on my view there is a fundamental sense in which it also exists “in the world,” since on my view undefined properties are constituents of the world that are independent of our linguistic conventions. Although he seems to think this is troubling, I don’t see that any problem has been demonstrated. In fact, I am willing to take the point a step further by maintaining the existence of vague objects – i.e. objects *o* which are such that for certain bits of matter *m* it is undefined whether *m* is part of *o*. Consider, for example, Mount Olympus, the highest mountain in the Olympic mountain range in western Washington state. For anyone familiar with the mountain, it seems evident that the boundary separating it from various nearby peaks is vague and indeterminate. In my view, this means that for certain masses *m* of rock the claim that *m* is part of Mount Olympus and the claim that *m* is not part of Mount Olympus are undefined. It follows from this – doesn’t it? – that there is something *x* in the world such that the claim that *m* is part of *x* is undefined. If this is not a clear case of undefinedness being “in the world, independent of us” then I don’t know what could be meant by the phrase.^{vii}

Reply to Tappenden

I now turn to Jamie Tappenden’s comments. As he points out, our views about vagueness and the Liar are similar. Being comrades in arms, it is not surprising that each of us finds himself to be in broad agreement with the other. Thus, I welcome his observations, and have no serious objections to them. I do, however, have two minor clarifications to offer. First, in characterizing his extension of my smidget example,

Tappenden stipulates that every member of a certain group A is a smidget, that every member of B is not a smidget, and that no decision is made about persons outside the two groups. He says “it is, of course, crucial that the smidgethood of the indeterminate cases is genuinely *left open*,” and he draws attention to cases in which language users find significant practical reasons to employ predicates that are only partially defined, while deferring until later the question of how initially undefined cases are to be settled, if at all. He notes that the ability to accommodate this fact in a natural way is an important advantage of accounts of vagueness of the sort we favor over rival accounts like epistemicism. All of this is correct and important. What I want to clarify is the sense in which speakers using a predicate that is initially undefined for certain objects leave it open whether the predicate is to apply to those objects. They leave it open in the sense that they retain the discretion, in the future, to slightly adjust the semantic content of the predicate (the property it expresses) so as to bring it about that it applies, or does not apply, to the objects in question. On the account of vagueness that Tappenden and I favor, it is a feature of vague predicates that they are both partial and context sensitive, and that their meanings include principles governing permissible changes in their semantic content from one context to another.

My second point of clarification involves what Tappenden calls *local consistency rules*, roughly “*sentences that cannot be assigned false on any increase in precision [of vague predicates], though they also currently receive no truth-value at all, if our account of the connectives is truth-functional.*” Tappenden’s thought is that such sentences constrain increases in the precision of vague predicates in the sense that the meanings of those predicates preclude any assignments of extension and antiextension that would

falsify the local consistency rules. He uses this idea to illuminate what is right about even pathological Tarski-biconditionals like *the Liar is true iff the Liar is not true*. What is right about them is that the meaning of the truth predicate is such that these biconditionals cannot be made false (though they also cannot be made true).^{viii}

Again, I have no quarrel with this basic idea. However, it is important to forestall a potential misunderstanding. To say of a local consistency rule *s* that *s* is presently undefined, but that the meaning of the vague predicate *s* contains is such that *s* cannot be made false, is **not** to say that *s* cannot be, and hence **is not**, false. Rather, it is to say that the rules presently governing the expressions in *s* are such that they do not determine any characterization of *s* as either true, false, or untrue, and that the meaning of the vague predicate in *s* is such that in any context the extension and antiextension determined for the predicate will be such that the rules governing *s* will not determine that *s* is false. A similar point holds about the content of the claim that *the Liar is true iff the Liar is not true* cannot be **made** true.

Reply to Williamson

I will conclude by replying to Timothy Williamson. Because he is an advocate of an opposing account of vagueness, it is not surprising that several of his comments are sharply critical. However, they also raise important issues that I hope my to illuminate.

Williamson's first main point concerns apparent problems with my use of strong Kleene conditionals. The point is illustrated by a variant of one of his examples.^{ix} We begin by assuming that the vague predicate *bald* is undefined for a certain man Joe, as well as for those with slightly less, or slightly more, hair. Next we assume that the

conditionals in (6) and (7) are strong Kleene versions of the material conditional, governed by table (2).

- 6a. If Joe is bald, then Joe is bald
- b. If Joe is bald, then he would be bald if he had one **less** hair
- c. If a person is bald, then he would be bald if he had one **less** hair.
- 7a. If Joe is bald, then he would be bald if he had one **more** hair ,
- b. If a person is bald, then he would be bald if he had one **more** hair.

On the account I have offered, all these conditionals are undefined, and so should be rejected. Intuitively, however, there is a difference between those in (6), which appear true and acceptable, and those in (7), which are clearly questionable. Certainly (7b) cannot be accepted or taken to be true, and in some contexts it would seem that (7a) shouldn't either. The theory owes us an explanation of this contrast.

There are two options. Either at least some of the conditionals should be given semantic interpretations in which they are **not** strong Kleene conditionals, or they should be shown to have uses in which they **convey** truths that differ from the propositions they semantically express. Two possibilities regarding the semantic option are: (i) to introduce a conditional like (1) that is true when both antecedent and consequent are undefined (and false when the antecedent is undefined but the consequent is false), and (ii) to introduce a special supervaluationist conditional that is true when every means of closing gaps results in a truth. Either way, the conditionals in (6) will be characterized as true, even though the inductive premise (7b) of the Sorites paradox will not.

According to the nonsemantic option, there is no need to move beyond the strong Kleene conditional. Let us focus on the contrast between (6c) and (7b). The conditional (6c) can be used to communicate the metalinguistic truth, (M6c).

M6c. No matter what contextual standards for baldness we adopt, if according to those standards ‘he is bald’ applies to x, then according to those standards ‘he is bald’ would apply to x if he had one **less** hair.

On this view, (6c) is semantically undefined, but the speaker may use it to convey the genuine truth (M6c).^x Since (7b) has no such use, we can account for the intuitive contrast nonsemantically. Thus Williamson is wrong when he suggests that the problem can’t be avoided at the metalinguistic level. It can, but doing so requires a metalinguistic formulation other than the one he discusses.

I prefer the nonsemantic option for two reasons. First, it is unclear to me, apart from questions involving vague or partially defined predicates, what the proper semantic analysis of indicative conditionals in English is. Hence in defending my account of vagueness, I would rather not put too much weight on one particular analysis. Second, it seems likely that the nonsemantic option will remain viable no matter what semantic analysis of conditionals ultimately proves correct for English. Moreover, problems of vagueness, including the Sorites paradox, can be made to arise in languages in which the conditionals are stipulated to be strong-Kleene. That, in fact, is the best way of looking at my chapter on vagueness. From that perspective, the contrast between the conditionals in (6) and (7) is explained nonsemantically.

Williamson’s next worry involves principle (P2*) of chapter 7.

P2*. For any two patches of color x and y that are perceptually indistinguishable to competent observers under normal conditions, if someone who is presented with x characterizes the predicate *looks green* as applying to it, then that person is thereby committed to a standard that counts the predicate as applying to y as well.

About this, Williamson says the following:

“Consider a case in which two patches of color x and y are perceptually indistinguishable to competent observers under normal conditions, and I am presented with x , but it is vague whether I have characterized ‘looks green’ as applying to x ; I assent to the application in a somewhat reluctant, hesitant or frivolous way. Thus, given the view expressed in (P2), it is also vague whether I am thereby committed to a standard that counts the predicate as applying to y as well. Thus, on Soames’s view, the conditional in (P2*) has an undefined antecedent and consequent in this instance. Therefore, (P2*) is not true on Soames’s semantics, contrary to his claim.”*

Although Williamson raises a significant issue here, in the end I don’t think it is decisive. I begin with a basic point. The fact that one assents to something hesitantly, reluctantly or frivolously does not show that it is vague that one has assented to it; and if one has assented to the claim that *looks green* applies to x , in whatever way, then one is committed to the standard indicated in (P2*), even if that commitment was made in a hesitant, reluctant, or frivolous way. However, let us suppose, for the sake of argument, that it is possible for it to be vague whether a person has assented to the relevant claim, and hence vague whether that person has (explicitly) characterized *looks green* as applying to x , no matter how this might come about.^{xi} Since, under this assumption, the

predicate *y* characterizes the predicate 'looks green' as applying to *x* is vague, it will follow that if the conditional in (P2*) obeys the strong Kleene table, then, (P2*) may turn out undefined. Even then, it doesn't follow that (P2*) **must** be undefined, since it is not precluded that the person offering (P2*) hasn't adopted a contextual standard that assigns to the vague predicate an extension and antiextension that exhausts the range of possible cases, -- in which case (P2*) will express a truth, after all. However, let us further stipulate that this is not so, and that (P2*) is undefined in the context in question. Despite this, it may still very well be that (P2*) **appears** true, and for that reason continues to play the role (described on pp. 214 - 215 of Understanding Truth) of making the inductive premise of a Sorites argument appear true, as well.^{xii} If so, then we face the additional task of describing the nearby truth that makes both appear true.

With this in mind, consider (P2**).

P2**. For any two patches of color *x* and *y* that are perceptually indistinguishable in color to competent observers under normal conditions, if someone who is presented with *x* clearly and determinately characterizes the predicate *looks green* as applying to it, then that person is thereby clearly and determinately committed to a standard that counts the predicate as applying to *y* as well.

Here, we seem to have eliminated worries about vagueness and undefinedness by prefixing a *clearly and determinately* operator to the relevant clauses. Under this scenario, we use the strong Kleene conditional to state the nearby truth (P2**) that makes both (P2*) and (P2) appear true.

Might someone continue to object that there are cases of higher-order vagueness that cannot be eliminated in this way? Perhaps, but by now we have gone pretty far down a

dubious road. To have gotten this far one would have to assume two things: (i) that the predicates *y assents to the claim that 'looks green' applies to x* and *y characterizes 'looks green' as applying to x* are vague – a position that makes sense only if one is willing to regard almost every predicate of nonmathematical objects as vague -- and (ii) that this vagueness can never be eliminated -- either by adopting jointly exhaustive extensions and antiextensions in particular contexts, or by using the *clearly and determinately* operator to create totally defined predicates when needed. In my view, the conjunction of (i) and (ii) is implausible. If we reject it, then Williamson's objection can be handled in the ways already indicated. If, however, one accepts it, then I suspect one **will** have to adopt some new conditional, either in addition to, or in place of, the strong Kleene conditional -- quite apart from the Sorites-type considerations we have been dealing with here. In this eventuality, the theorist may well want to introduce the conditional operator (1) which yields determinate truth-values even in cases in which the arguments to it are undefined. When (P2*) is construed as containing this operator, it is true, even though the inductive premise (P2) of the Sorites argument still must be rejected.^{xiii} Even if one does reach this position, nothing in it undermines the basic analysis of vague predicates as both context sensitive and partially defined.

Williamson's third point concerns the relationship between context sensitivity and partial definition. On my analysis, vague predicates are both context sensitive and partially defined. Williamson appears to think that a case can be made for context sensitivity, but that this case fits better with a treatment of vague predicates as totally defined than it does with my analysis of them as partially defined. Among his comments are the following:

“Soames’s account is based on the old idea that borderline cases involve truth-value gaps. He says rather little in justification of this idea, merely that it seems arbitrary where to locate the cut-off point for a vague term and that it is unclear what semantic mechanism could determine that point.”

*“Soames uses it [the mechanism of contextual adjustment of boundary lines] to explain how sharp semantic lines can be systematically elusive and therefore appear not to exist. **But if he is willing to do that in defense of a hidden sharp boundary between truth and undefinedness, would it not be simpler to do it [i.e. use the contextual adjustment of boundary lines] in defense of a hidden sharp boundary between truth and falsity?. ...Soames can hardly take his generic objections [arbitrariness and unclarity of semantic mechanisms] to truth values in borderline cases to be decisive, for they could also be made to the claim that a predicate is either true, false or undefined in a higher-order borderline case.**” (my emphasis)*

*“Some of the conversational mechanisms postulated by Soames may well have a part to play in a full account of vagueness in natural languages. **But contextual variation in reference is entirely consistent with epistemicism about vagueness.**” (my emphasis)*

My view is just the opposite of Williamson’s – once one recognizes the context sensitivity of vague predicates, the conclusion that they are also partially defined is well nigh irresistible.

To say that vague predicates are context sensitive is to say that they are indexical. While the semantic content of an indexical varies from one context of utterance to

another, its meaning does not. Rather, its context-invariant meaning constrains the indexical to take on semantic contents with certain specified features. Sometimes these constraints identify semantic content in terms of a fixed contextual parameters – e.g. the content of ‘I’ is the agent of the context, the content of ‘now’ is the time of the context, and the content of ‘actually’ is the world of the context. In other cases, the meaning of an indexical constrains its semantic content to be one that satisfies a certain condition – e.g. the content of ‘he’ must be male, the content of ‘she’ must be female, and the content of ‘we’ must be a group of individuals that includes the agent of the context. A speaker using one of these indexicals is free to select any salient content that satisfies the relevant constraints.

If, as I believe, vague predicates are context sensitive, then this is the model on which they must be understood. Thus, if the predicate *looks green* is context sensitive, then it has a context-invariant meaning that constrains the properties it semantically expresses in different contexts. There is little alternative but to take these constraints as involving what I call the predicate’s *default determinate-extension* and *default determinate-antiextension* – these being, respectively, the set of things that the communitywide rules of the language (plus relevant nonlinguistic facts) determine that the predicate applies to, and the set of things that the communitywide rules of the language (plus the relevant nonlinguistic facts) determine that the predicate does not apply to (in any context in which the predicate is successfully used with its normal literal meaning).^{xiv} Given these exemplars of *looking green* and *not looking green*, we may view the meaning of the predicate as constraining the property it semantically expresses in a context to be one that satisfies roughly the following condition.^{xv}

Condition on the Semantic Content of the Predicate *Looks Green* in a Context

For any context *c*, if *looks green* semantically expresses a property *P* in *c*, then the members of the default determinate-extension of the predicate possess *P* and the members of the default determinate-antiextension do not.

In any given context, speakers have the discretion of selecting a contextually salient and relevant property from the range of properties satisfying the consequent of this condition.

Sometimes the property selected will be one that is possessed by some objects over and above those in the default determinate-extension, and one that is not possessed by other objects over and above those in the default determinate-antiextension. Which objects these are in any given case depends on the interests and purposes of the conversational participants .

On this picture, the default determinate-extension and the default determinate-antiextension play crucial roles in establishing two distinctions that are fundamental to the semantics of context-sensitive, vague predicates. First, they are used to distinguish cases in which the speaker asserts something that is definitely false (untrue) from cases in which the speaker's remark is properly regarded as true due to the legitimate exercise of discretion in setting the context-sensitive standards governing the predicate. For example, when speaker explicitly characterizes an object *o* of a certain sort as "looking green", or as "not looking green", he may, depending on other relevant features of the conversation, succeed in adjusting the boundaries of the predicate so as to make his remark true. However, if *o* falls outside a certain range, the speaker's remark can only be regarded as false (provided that he is using the predicate with its normal literal meaning). The default determinate-extension and the default determinate-antiextension determine this range. Second, these notions are needed to distinguish the context-invariant meanings of context-sensitive vague predicates of the same type. For example, we distinguish the meaning *looks green* from that of *looks yellow* by citing

their different default determinate-extensions and antiextensions. If these predicates were not semantically associated with such sets, then each could correctly be applied to items of virtually any appearance, and we would be hard-pressed to indicate how they differ in meaning.

Since these points apply to any theory of context-sensitive vague predicates, any such theory must countenance default determinate-extensions and default determinate-antiextensions. The only remaining question is whether this distinction is best incorporated in a theory that analyzes typical vague predicates as partially defined, or one that requires them to be totally defined. The decisive factor in answering this question involves the lines that must be drawn by the two theories over and above those that define default determinate-extensions and default determinate-antiextensions. If vague predicates are partially defined, then these extra lines are relatively easy and straightforward to locate. However, if vague predicates are required to be totally defined, then the extra lines end up being mysterious and arbitrary.

To see this, first consider an ordinary conversation between two speakers who make no special assumptions in advance about what looks green and what does not. Both are competent speakers of English who understand the predicate, but prior to this point in the conversation there has been no discussion of the colors things appear to have, and no special standards have been adopted. They now notice a car approaching from a distance, and one of the speakers, who can't make out the car clearly, asks the other what color it is. The second responds by saying "It looks green." Suppose that the car does look green, and that its appearance places it well within the default determinate-extension of the predicate. What property does the speaker attribute to the car, and the predicate semantically express, in this context? According to the theory that treats vague predicates as partially defined, the answer is clear. The property is one

that applies to all members of the default determinate-extension, does not apply to all members of the default determinate-antiextension, and is undefined otherwise. There is nothing mysterious or arbitrary about this; the property is the default semantic content of the predicate. As such it is expressed in any context in which speakers make no special assumptions about the range of cases over which they have discretion.

Things are not so simple for the theory that treats vague predicates as both context sensitive and totally defined. No matter what the context, the theory constrains the predicate to express a property that applies to members of the default determine-extension and does not apply to members of the default determinate-antiextension. However, the theory also requires the property expressed to be one of the many totally defined properties that satisfy this condition. The problem is that we have no principled basis for making a unique choice among these different properties; any such choice seems arbitrary, and it is mysterious what semantic mechanism one might use to make it.

A similar problem exists in cases in which speakers exercise their discretion by explicitly characterizing an object *o* as looking green, where *o* lies in the intermediate range between the default determinate-extension and the default determinate-antiextension. In cases like this the theory of partially defined predicates analyzes the predicate as expressing a property *P* that is possessed by *o*, everything perceptibly indistinguishable in color from *o*, plus everything that “looks greener than *o*” in the sense of looking more similar in color to items in the default determinate-extension than *o* does. Typically, *P* will not be possessed by anything that was not in the determinate-antiextension of the predicate according to the conversational standards prior to the new characterization of *o*, and *P* will be undefined for everything else.^{xvi} As before, the choice of this property is not arbitrary. The semantic content of the predicate in

this context is a property the range of application and nonapplication of which represents the minimal change from previously existing conversational standards required by speakers' explicit decision to characterize the predicate as applying to *o*.

As before, things are more problematic if vague predicates are required to be totally defined. We know that when *o* is in the intermediate range of the predicate speakers have the discretion to characterize *o* either as looking green or as not looking green. We also know that if they explicitly characterize *o* as looking green, then the predicate will come to express a property that determines an extension that includes everything perceptibly indistinguishable in color from *o*, plus everything that "looks greener" than *o*. But if we insist that the property be defined for all objects, it remains an open question what the antiextension will be. Though many answers are conceivable, there seems to be no principled basis for selecting one over all the others. One simply cannot tell merely from the agreement to count *o* as looking green what other objects are now to be counted as not looking green. Once again, the theory of totally defined, context-sensitive vague predicates forces an arbitrary choice on us while leaving it mysterious what semantic mechanism could be used to make that choice.

These results show that if vague predicates are context-sensitive, then they are also partially defined. Hence, in my opinion, Williamson is wrong in thinking that the benefits of context sensitivity can be had without partial definition. Context sensitivity and partial definition are a package deal.

This still leaves the task of giving a plausible, nonarbitrary explanation of how the default determinate-extensions and antiextensions of vague predicates are determined. This problem is genuinely troublesome. While I don't have any final resolution to offer, I am guided by the following picture. Let *E* be the set of those things such that every speaker of the

language who understands the predicate, and who has normal eyesight, would judge them to look green when placed in appropriate conditions. Let A be the set of things such that every such speaker would judge them not to look green. E and A are candidates for being the default determinate-extension and default determinate-antiextension of the predicate. Everything else is in the undefined range, within which speakers are free to adjust the boundaries of the predicate. The idea is that to characterize something in E as not looking green would be to give evidence that one doesn't understand the predicate (provided that there is nothing wrong with one's eyesight or the conditions in which one viewed the object). However, an item that is a member of neither E nor A – i.e. an item such that some competent speakers would be inclined, when viewing it in good conditions, to characterize it as looking green and others would not be so inclined, or would be inclined to make the opposite determination -- is an item for which no clear authority can be cited to resolve apparent disagreement. (It is not as if we have a standard sample, on analogy with the standard meter, which precisely fixes the extension of the term.) In such a case it is natural to regard the object as lying within the realm of speakers' discretion.

On this picture, it is not arbitrary what the default determinate-extension is: it is the set of things about which there is essentially no disagreement among people who both understand the expression and who are competent observers. Moreover, we have some idea how linguistic practice could determine where this line is drawn -- it is drawn so as to include only those objects the characterization of which as not looking green would be regarded by normal speakers as evidence that the person did not understand the predicate, or was perceptually deviant. The problem with this picture is not, I think, its plausibility. Rather, what gives one pause is its reliance on the background concepts of being a competent speaker, of

understanding the meaning of a word, of being a normal observer, and of being a piece of linguistic behavior which provides evidence that one either does not understand the predicate or is perceptually deviant. In order to complete the picture, we would need to give precise, noncircular explications of all these concepts – a task made more problematic by the suspicion that they may themselves be vague. In my view this the most serious unsolved problem confronting the theory.

Williamson's final major objection to my view concerns the contents of vague predicates, the linguistic conventions governing their use, and the role of specific individuals as exemplars in determining their extensions and antiextensions. He quotes the following passage from page 210 of Understanding Truth regarding how the default determinate extension and antiextension of *bald* are determined.

“(i) The predicate is applied to certain specific individuals who are taken as paradigmatic examples of baldness, and it is denied of other specific individuals who are taken as paradigmatic examples of people who are not bald. (ii) It is resolved that other individuals, y , are to be taken as determinate instances of baldness if they bear the relation $R_e - y$ has a perceptually equivalent amount of hair as x or less hair than x – to someone, x , in the initial sample of paradigmatically bald individuals. (iii) It is further resolved that still others, z , are to be taken as determinate instances of individuals who are not bald if they bear the related relation $R_a - z$ has a perceptually equivalent amount of hair as w or more hair than w – to someone, w in the initial sample of individuals who are paradigmatically not bald.”

About this Williamson says:

“Evidently this is a fantasy. There are no specific individuals who are semantically guaranteed to be bald. Although we can point to certain specific individuals as paradigms of baldness, no convention of English implies that they are not in the antiextension of ‘bald’ and wearing a tight skin-colored cap... If there are semantic conventions of English about ‘bald’, their content is general; they make no reference to specific individuals”.

I am puzzled by this. In my view, the semantic content of *bald* (in a context) is a general property that does not essentially involve any individual, and there is no individual whose existence or baldness is constitutive of the meaning of *bald* in English. So it is not clear to me precisely what Williamson is objecting to. He says that the semantic conventions of English are general in their content, and make no reference to specific individuals. There may well be a sense in which that is true. But there is also a way in which individuals can play an important role in determining the content of general terms.

Forget, for the moment, about vagueness. Think of natural kind terms that are introduced ostensively with reference to individuals in a certain sample. For example, we might say, *Let the predicate K apply to all members of the species of which most of them* (demonstrating particular individuals) *are instances*. This stipulation referring to particular individuals may be the mechanism by which K acquires its content – even though that content is not the same as the content of the expression used to introduce the predicate, and even though there is no individual whose existence and membership in a

certain species is constitutive of the meaning of K in the language. Moreover, K could have acquired the same content without any ceremony or explicit stipulation, simply by being used repeatedly as a predicate of certain specific individuals, with the understanding that it is to apply to all and only members of the species that (most of) those individuals belong to.^{xvii}

There is no evident reason why something similar might not hold for many vague predicates. For example, the default determinate extension of *bald* might be determined by the stipulation *A person is bald if he has that much hair (or less) on his head*, or *A person is bald if he has the same amount of hair, or less, on his head as most of them (appear to) have* said while demonstrating certain individuals, even though the content of antecedent clause of the stipulation is not included in the content of *bald*.^{xviii} More generally, the same result could be achieved without explicit stipulation. It may either be stipulated, or become clear through ordinary use, that the predicate is to apply to this, that, and the other individual, taken as paradigm cases of baldness, while not applying to these or those different individuals, taken as paradigm cases of nonbaldness, while also being clear that it applies (doesn't apply) to any other objects that bear certain relations to the individuals taken as paradigmatic cases.

Finally, Williamson adds a comment about what he takes to be the difficulty of stating the general linguistic conventions governing vague predicates on my account.

“For, in view of higher-order vagueness, the conventions governing vague terms will themselves be vague. Since the homophonic approach to stating the conventions fails, for the reason noted above [namely, context sensitivity], a semantic convention about a vague term t will need to be

stated by use of another vague term t. Soames does not discuss the danger of an infinite regress. Nor would it be attractive to postulate unstatable conventions, for the most powerful accounts of conventions make them a matter of mutual knowledge and belief.” [my addition]*

Perhaps Williamson says this because he takes higher-order vagueness to be ubiquitous and pervasive throughout language. For myself, I believe that the linguistic conventions governing vague predicates can be nonhomophonically stated, just as corresponding conventions for other context sensitive terms can. For example, the convention governing the reference of the demonstrative ‘he’ is (roughly) that it refers to the salient male that is demonstrated in the context. Note, the convention does not specify how one determines which male is demonstrated in the context, nor does the convention guarantee that there never is any vagueness regarding which individual is demonstrated. Similarly, the convention governing the determinate-extension of a vague predicate F in a context is that it includes the default determinate-extension of F plus all individuals that bear a certain contextually determined relation R to things that have been explicitly characterized as “being F” in the context. In many cases, the meaning of F will put constraints on the range of relations that may qualify as the relevant R in any context – while leaving open the choice among these relations in particular contexts. Even if it is sometimes vague which relation plays this role in a given context, it is not clear why this is any more of a problem with vague predicates than it is with ‘he’.

ⁱ Here, and throughout, I use boldface italics to play the role of corner quotes.

ⁱⁱ This leads me to think that Gupta may have something else in mind. Perhaps he thinks that there is another sense of *if and only if* -- different from the one I defined and he repeated -- such that a language can contain a biconditional operator with that sense along with its own truth predicate. That is something I agree with, as I indicated in the book. But it remains to be seen how this constitutes a criticism of the model I presented.

ⁱⁱⁱ In this paragraph I follow McGrath in using ‘p’ as a schematic sentence letter, rather than as a variable over propositions.

^{iv} Given the two principles McGrath actually states (one about the correctness of belief and one about the correctness of disbelief), plus (i) the assumption $p \text{ or } \sim p$, and (ii) the principle *if it is correct to believe p, then it is not correct to disbelieve p*, one can derive the conclusion *it is not correct to disbelieve both p and $\sim p$* (i.e. one of the disbeliefs must be incorrect). But whatever one thinks about (ii), it is a central part of my view that (i) cannot be accepted in the cases under discussion.

^v ‘*’ in the tables is not a third truth-value, but rather means ‘undefined’.

^{vi} Assuming that the new facts do not alter one’s initial judgement about the actual height (or age) of o.

^{vii} Although in the book I talk about vagueness only in connection with predicates, I also have a view about names. For example, I believe that there is a certain object x (Mount Olympus) such that the name *Mount Olympus* determinately refers to x, and x is not identical with any object y the boundary conditions of which are fully precise.

^{viii} Presumably, Tappenden here tacitly takes *iff* to be a Kleene-style connective, and not my “just in case” biconditional operator \equiv .

^{ix} This example, and the criticism based on it, was also put forward by Hartry Field in his contribution to the “Author Meets Critics” session on Understanding Truth at the Pacific Division Meetings of the American Philosophical Association in Albuquerque, New Mexico on April 6, 2000. My reply to Williamson on this point follows the remarks I made there.

^x To say that according to the standards in the context *he is bald* applies to *x* is to say that the contextual standards plus the world of the context together determine that *he is bald* applies to *x*.

^{xi} If *assents to the claim that ‘looks green’ applies to x* is vague, is every predicate vague? If every predicate is vague, how do we make sense of the intuitive contrast between that which is vague and that which is not?

^{xii} (P2) is as follows: For any two patches of color *x* and *y* that are perceptually indistinguishable in color to competent observers in good light under normal conditions, *x* and *y* look to be the same color and so one looks green if the other one looks green.

^{xiii} There is a potential worry about (P2*) that Williamson does not raise which nevertheless should be addressed. Suppose we have a Sorites series in which the boundary line between the cases for which the predicate *looks green* is undefined and the members of its default determinate-antiextension falls between x_i and x_{i+1} – the former being that last member of the series for which the predicate is undefined and the latter being the first member of the series in its default determinate-antiextension. Suppose further that a subject who is presented with x_i decides to characterize it as looking green. Three questions need to be answered. (i) What has the subject asserted by virtue of saying of x_i *That looks green?*, (ii) What property is semantically expressed by the predicate in this context?, and (iii) What standard is the subject committed to by virtue of characterizing *looks green* as applying to x_i ?

For other undefined x_n , earlier than x_i in the series, answering these questions is straightforward.

(i) When the speaker says of x_n *That looks green*, he asserts a proposition that attributes to x_n a certain property P_n which applies to all objects that look greener than x_n , as well as to x_n plus all objects that are perceptually indistinguishable in color from x_n (in a certain specific sense defined in my chapter on vagueness). Among these objects is x_{n+1} , the next item in the series. (ii) Since *looks green* is indexical, P_n is the property semantically expressed by the predicate in the context. (iii) The speaker is committed to a standard that counts x_{n+1} as being on a par with x_n in two respects: (a) the property P_n that the speaker has asserted x_n to have is a property that x_n shares with x_{n+1} , and (b) as the speaker is using the predicate in the context, the formula *y looks green* semantically expresses a truth no matter whether ‘y’ is assigned x_n or x_{n+1} as value.

The situation is more complex when the speaker characterizes the predicate as applying to x_i . When, in context c , the speaker says of x_i *That looks green*, the rule for extending the predicate to previously undefined cases would seem to apply, with the result that the speaker asserts a proposition that attributes to x_i a property P_i which applies to all objects that look greener than x_i , as well as to x_i plus all objects that are perceptually indistinguishable in color from x_i (including x_{i+1} , the next item in the series). This is just as before. However, since x_{i+1} is in the default determinate antiextension of the predicate, **no** context in which the predicate is used with its normal literal meaning can be one in which the property it semantically expresses applies to x_{i+1} . Hence, P_i can’t be the property semantically expressed by *looks green* in c .

What should we say about this case? In my view, the two most promising alternatives are as follows: (A1) No property is semantically expressed by the predicate in this context, owing to the fact that two semantic constraints on the predicate conflict. The rule governing the extension of the predicate to new cases constrains its semantic content in c to be P_i , whereas the rule governing the default determinate-antiextension prevents it from being P_i . (For any context in which the predicate is used with its normal literal meaning to semantically express a property P , it cannot be the case that P applies to items in the default determinate-extension of the predicate.) Since no property is semantically expressed, when the speaker says of x_i *That looks green*, the sentence he utters fails to semantically express any proposition.

Nevertheless, he does seem to assert something; in particular he asserts the proposition that attributes P_1 to x_i . In short, everything is as before, except that – because of the conflict -- this proposition does not count as the semantic content of the sentence in c . If this is right, then the speaker is committed to a standard that counts x_{i+1} as being on a par with x_i in the sense of (iiia) above, but not (iiib). As a result (P2*) can be retained in essentially its present form. (A2) On this alternative, the conflict between the two semantic constraints is resolved by selecting as the property semantically expressed by the predicate in c the nearest available semantic content, P_{i-1} , that would escape such conflict. This property applies to all objects that look greener than x_{i-1} , as well as to x_{i-1} plus all objects that are perceptually indistinguishable in color from x_{i-1} , including x_i . Hence, the speaker's use of *That looks green*, said referring to x_i , semantically expresses the proposition that attributes P_{i-1} to x_i . This proposition is both true and asserted by the speaker. On this view, the speaker is **not**, strictly speaking, committed to a standard that counts x_{i+1} as being on a par with x_i , with the result that (P2*) is not literally true. At best one can think of it as a sort of default rule that applies except when forced to give way by independent semantic constraints. Although this alternative requires amending my discussion in Understanding Truth, it would not, in my opinion, drastically change the role of (P2*) in the overall argument. If speakers confuse the Sorites premise (P2) with a genuine default rule encoded in the meaning of the predicate, then it is understandable why the Sorites premise is apt to seem more plausible than it really is.

This note responds to issues raised by Teresa Robertson in “On Soames’s Solution to the Sorites Paradox,” Analysis, 60:4, 2000, 328-334.

^{xiv} Understanding Truth, p. 209.

^{xv} Three further points are important to keep in mind in understanding this condition. (i) It is not the only constraint that the meaning of *looks green* places on its semantic content in a context. For example, suppose x and y are members of neither the default determinate-extension nor the default determinate anti-extension of the predicate. Suppose further that x “looks greener” than y , in the sense of looking more similar in color to the members of the default determinate-extension of the predicate than y does. Then for

any context c and property P , if P is the semantic content of the predicate in c , then y possesses P only if x does. (ii) In stating both (i) and the condition in the text, I mean to be making claims that that are not themselves subject to being undefined. As my discussion of Williamson's second objection indicates, there are a variety of ways in which this result can be achieved. For present purposes it is sufficient to interpret talk of an object possessing a property P , or not possessing P , as talk of the object's determinately possessing P , or determinately not possessing P . (Similarly for talk of P 's being semantically expressed by the predicate in the context.) (iii) In saying that the meaning of *looks green* constrains its semantic content to satisfy the condition given in the text, I am saying that the meaning has this effect, not that the condition is something that is explicitly known by speakers. Speakers do not have to explicitly possess the concepts of the default determinate-extension and the default determinate-antiextension of the predicate, nor need they have any explicit *de re* beliefs about these particular sets. Rather, an agent who understands the predicate typically is acquainted both with some members of the default determinate-extension, which the agent treats as exemplars of *looking green*, and with some members of the default determinate-antiextension, which the agent takes as exemplars as *not looking green*. The agent takes the former to look a certain way – i.e. to have a certain property P_{DDE} – and the latter to look a different way – i.e. to have a different property P_{DDA} . In addition, the agent knows that although what properly counts as “looking green” or “not looking green” may vary somewhat depending on the context, things that look one way – i.e. have the property P_{DDE} -- always count as “looking green”, while things that look the other way – i.e. have the property P_{DDa} -- always count as “not looking green”.

^{xvi} There are some cases that are more complicated than this, but they are relatively rare, and don't, I believe, substantially affect the main point here.

^{xvii} For a more detailed explanation of the role of ostensive definition in determining the contents of many natural kind predicates see chapter 10 of my *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*, (New York: Oxford University Press), 2002.

^{xviii} As indicated in point (iii) of note 15, the role of positive exemplars of a vague predicate is typically to acquaint speakers with a property P_{DDE} possessed by those exemplars that determines the default determinate-extension of the predicate.