formalized languages, appeared in A. Tarski 1956, *Logic, Semantics and Metamathematics: Papers by Alfred Tarski 1923–1938*. Oxford: Clarendon Press.
Tarski, A. 1944. The semantic conception of truth and the foundations of semantics. *Philosophy and Phenomenological Research* 4: 342–60. Reprinted in *Truth*, ed. S. Blackburn and K. Simmons. 1999. Oxford: Oxford University Press.

On Soames's solution to the sorites paradox

TERESA ROBERTSON

Scott Soames (1999, ch. 7) has recently offered a new solution to the sorites paradox. Although this solution has some appeal, it seems to me that, short of some substantial revision, it fails.

1. *Presentation of Soames's solution*

Soames's solution to the sorites paradox turns on two features of vague predicates: (i) they are (at least potentially) partially defined and (ii) they are context-sensitive. A partially defined predicate, 'is *F*' say, is one whose extension and antiextension are mutually exclusive but not jointly exhaustive. For any object *o* that is in neither the extension nor the antiextension, both the claim that *o* is *F* and the claim that *o* is not *F* should be rejected. There is (at least potentially) a truth-value gap.[1] Given the understanding of the material conditional that is provided by the strong Kleene tables,[2] this means that for any standard sorites paradox of the form of the (implausibly short) one displayed in §2, the conditionals near the beginning will be true; the ones in the middle will lack truth values – that is, will be 'undefined'; and the ones toward the end will be true, since their antecedents will be false. None will be outright false. So the conclusions of sorites arguments can be avoided by rejecting some of the conditional premisses without thereby being forced to accept their negations.

By itself, this is not very satisfactory. Soames considers two objections.

Objection 1. If partial definition were the whole story and *o* were an object in the undefined range of the predicate, then we would have no

[1] The gap would be merely potential in the case of 'is bald' if, for example, all people had no hair.

[2] Where '*U*' is used for neither *T* nor *F*, we have the following: (*T*, *T*) yields *T*; (*T*, *U*) yields *U*; (*T*, *F*) yields *F*; (*U*, *T*) yields *T*; (*U*, *U*) yields *U*; (*U*, *F*) yields *U*; (*F*, *T*) yields *T*; (*F*, *U*) yields *T*; (*F*, *F*) yields *T*.

ANALYSIS 60.4, October 2000, pp. 328–34.

choice about how to characterize *o*. We would have to reject the claim expressed by '*o* is *F*' as well as the claim expressed by '*o* is not *F*'. But it is commonly thought that speakers have some discretion in their use of vague predicates: we can imagine situations in which it is perfectly correct to count a borderline case as a member of either the extension or the antiextension. (For example, we might acceptably say, 'For the purposes of this discussion, we'll count Yoko – *sotto voce*: who is really a borderline case of being rich – as rich.')

Objection 2. If partial definition of vague predicates were the whole story, then there should be a pair of objects, o_i and o_{i+1}, that are adjacent in the sorites series and are such that, if they were presented to us, we would – to speak truly – have to characterize them by saying of o_i 'it is *F*' while rejecting the corresponding claim about o_{i+1}. But it seems there are no such objects.

Soames assuages these worries by appeal to the context-sensitivity of vague predicates. Before I explain this appeal, I should point out one type of context-sensitivity that is not at issue. Predicates like 'is tall' admit of a kind of context-sensitivity that is not peculiar to vague predicates. In some contexts 'is tall' means *is tall for a former U.S. Presidential hopeful*; in others it means *is tall for a former professional basketball player*. 'Bill Bradley is tall' may express something true in the first context and something undefined in the second. But this has nothing to do with the vagueness of 'is tall'. This sort of context-sensitivity exists even with precise predicates such as 'has more children than average', which in some contexts means *has more children than the average American* and in others means *has more children than the average philosopher*. I will assume then that the context is sufficiently specified for this kind of context-sensitivity not to be at issue.

The context-sensitivity Soames is concerned with is that which allows speakers the discretion to adjust the extension or antiextension of a vague predicate by including initially undefined objects in the contextually-determined extension or antiextension – the kind of discretion exercised by saying, 'For present purposes, let's count Bill Bradley as tall for a former pro-basketball player'. It is easiest now simply to (nearly) quote Soames 1999: 209).

The Model. Vague predicates have the following characteristics.

(i) They are partially defined.

(ii) They have *default extensions and antiextensions*. The default extension of 'is *F*' is the set of things that the communitywide rules or conventions of the language (plus relevant nonlinguistic facts) determine that the predicate applies to. (Similarly for the default antiextension.)

(iii) Speakers have the discretion of adjusting the extension and antiextension of a vague predicate by including initially undefined cases – objects not in the default extension or antiextension in the predicate's *contextually determined extension or antiextension.*

(iv) Typically, this discretion is exercised by *explicitly characterizing* an object for which the predicate is initially undefined. When such an object *o* is explicitly characterized by a speaker as being *F* and other conversational participants accept this, the extension of the predicate is conversationally adjusted so as to include *o* plus all objects that bear a certain similarity relation R_e to *o*. (Similarly for explicitly characterizing *o* as not being *F*.)

(v) The meaning of the predicate constrains the similarity relation R_e. [A reasonable candidate for the similarity relation in the case of 'is tall' is this: yRx iff *y*'s height is greater than or equal to *x*'s height *minus* one millimeter.] (Similarly for R_{ae}.)

It is clear how this answers the first objection: clause (iii) allows for speakers' discretion. As for the second, that requires a longer story, which is best told by looking at what Soames calls the *dynamic version* of the sorites.

In this version we imagine a speaker being presented with items in the sorites progression one by one and being asked to characterize them. For each item she can either accept 'it is *F*', accept 'it is not *F*', or reject both claims. Soames's model is supposed to give rise to rules that govern the contextual adjustment of vague predicates. Such a rule would look like this.

> *Adjustment Rule.* If in a context C something *x* is explicitly judged to be *F*, then the extension of 'is *F*' in C includes everything bearing R_e to *x*. If in C something *y* is [explicitly] judged not to be *F*, then the antiextension of 'is *F*' in C includes everything that bears R_{ae} to *y*.[3]

We begin in the context according to which the extension and antiextension of 'is *F*' are identical to the default extension and antiextension respectively. The subject is presented with item 1, which she correctly asserts to be *F*. (For the rest of the paragraph, I virtually quote, selectively, from Soames 1999: 212–13.) The adjustment rule will determine that, by current conversational standards, item 2, which bears R_e to item 1, is also *F*. Since all the speaker is being asked to do is to characterize explicitly something as *F* that already counts as being *F* by previously accepted standards, fidelity to those standards dictates that she agree that item 2 is *F*. But once she has perceived item 2 and explicitly characterized it as *F*, the rule will classify item 3 as being *F*. After repeating this process a number of times, we might start at the other end with item *n*, which the subject correctly

[3] This rule is modelled on the rule Soames (1999: 212) gives for 'looks green'.

asserts not to be *F*. The rule will determine that, by current standards, item $n-1$ is not *F* either. Thus when we present item $n-1$ to the subject, fidelity to those standards will dictate that she agree that item $n-1$ is not *F*. But once she has explicitly agreed to this, the rule will classify item $n-2$ as not *F*. By continuing the process long enough, we may arrive at a point at which the subject characterizes some item *i* as not being *F* that previously she characterized as being *F*. Is this paradoxical? No. According to the analysis, when item *i* is initially characterized as being *F*, this is done with respect to a certain set of standards. Later, when it is characterized as not being *F*, this is done with respect to a different set of standards. There is no contradiction in the observation that something may be *F* with respect to one set of standards while not being *F* with respect to another.

If one considers carefully this account of the dynamic sorites, it becomes clear that some modifications to both Soames's semantic model and his adjustment rule are needed. To see this, it is helpful first to note that the '*x*' and '*y*' in the adjustment rule must *not* be assumed to range only over the default borderline cases. When the subject is presented with the last item in the default extension – which item is not of course a default borderline case – and judges it to be *F*, the rule for contextual adjustment presumably comes into play – which it would not if '*x*' and '*y*' ranged only over the default borderline cases – and allows that our subject thereby creates a context in which the next item is *F*. It seems simplest to assume then that the variables range over all the items in the sorites series.

This recognition of the variables' range makes it clear that Soames's semantic model is not capable of generating the adjustment rule. Clause (iv) says what happens only when a *default borderline case* is characterized as being *F* or not. Clause (iv), it seems, should be supplemented with something like the following.

> *Supplement to (iv).* When an object *o* in the default extension of 'is *F*' is explicitly characterized as being *F*, then if there are any objects *o'* that are not in the extension of 'is *F*', but that nevertheless bear R_e to *o*, the extension is conversationally adjusted so as to include those objects *o'*. (And similarly for the antiextension.)[4]

This recognition of the variables' range should also make clear that some modification to the adjustment rule itself is needed. Consider what happens if someone – who, along with her friends, has had one drink too many – judges that Shorty (who is in the default *antiextension* of 'is tall') is tall. According to the sort of rule in place, the speaker thereby creates a context in which Supershorty (who is just a little shorter than Shorty) is in fact tall.

[4] In correspondence Soames has endorsed this sort of supplementation.

Clearly this is not Soames's intent. Using the term *inextension* to denote the set of borderline cases, the rule should, it seems, be as follows.

> *Modified adjustment rule.* If in a context C something x, which is either in the default extension or default inextension of 'is F', is explicitly judged to be F, then the extension of 'is F' in C includes everything bearing R_e to x. If in C something y, which is either in the default antiextension or default inextension of 'is F', is explicitly judged not to be F, then the antiextension of 'is F' in C includes everything that bears R_{ae} to y.

To return to the main line now, Soames thinks that his treatment of the dynamic version of the sorites addresses the second objection, since for any of the conditional premisses in a sorites argument there is some context in which the conditional is true. Whenever o_i is presented to us and explicitly characterized as F, either this is false or a context is created in which o_{i+1} is also F. So although in every single context a 'counterexample' to at least one of the conditionals does exist (a 'counterexample' is a case in which the conditional is undefined or false), we can never *display* such a 'counterexample': as soon as we explicitly say, for example, 'o_i is F', if we speak the truth we create a context in which it is not the case that 'o_{i+1} is F' is undefined.

2. *An objection to Soames's solution*

To expose a problem with Soames's solution, I will discuss in detail a dynamic unidirectional sorites. To keep the details from overwhelming us, let's imagine that the default extension/inextension/antiextension for 'is F' is {1, 2}/{3, 4}/{5, 6}, where '1' names the first item in the progression and similarly for the other numerals. y bears R_e to x iff the number of y's position in the progression is less than or equal to one more than the number of x's position in the progression. (So 2 and 1 bear R_e to 1; 3, 2, and 1 bear R_e to 2 and so on.)

Given the default 'meaning' for 'is F', there are six possible contextually-determined 'meanings' for it.

context	*extension*	*inextension*	*antiextension*
U	1,2	empty	3,4,5,6
V	1,2,3	empty	4,5,6
W	1,2,3,4	empty	5,6
X	1,2	3	4,5,6
Y	1,2,3	4	5,6
Z	1,2	3,4	5,6

Here is the sorites argument.

(1) 1 is F
(2) If 1 is F then 2 is F

So (3) 2 is *F*
(4) If 2 is *F* then 3 is *F*
So (5) 3 is *F*
(6) If 3 is *F* then 4 is *F*
So (7) 4 is *F*
(8) If 4 is *F* then 5 is *F*
So (9) 5 is *F*
(10) If 5 is *F* then 6 is *F*
So (11) 6 is *F*

I will assume that context *Z*, which we can also call 'context 1', is the initial context. We present our subject with item 1. She explicitly characterizes it as being *F*. This need not induce any change of context. Indeed, assuming that context change is as conservative as possible, this will not induce any change. So, in context 1, sentence (1) is true. What is the status of sentence (2) in context 1? It is also true. (Just check the chart and the Kleene table for the material conditional.) Therefore, in context 1 sentence (3) is true. Now we present our subject with item 2. And she knows that in context 1, sentence (3), which says that item 2 is *F*, is true. She desires to be faithful to the standards of context 1, so she explicitly endorses sentence (3). When she does so, she changes the context: the modified adjustment rule demands that item 3 is in the extension of 'is *F*' in any context in which item 2 has been explicitly characterized as being *F*. This new context cannot be context 1, since in that context item 3 is not in the extension of 'is *F*'. Let's call the new context 'context 2'. Assuming that context change is conservative, context 2 is context *Y*. Sentence (3) is, of course, true in context 2. So is sentence (4). So sentence (5) is true in context 2. Now we present our subject with item 3. And she knows that in context 2, sentence (5), which says that item 3 is *F*, is true. She desires to be faithful to the standards of context 2, so she explicitly endorses sentence (5). When she does so, she changes the context: the rule demands that item 4 is in the extension of 'is *F*' in any context in which item 3 has been explicitly characterized as being *F*. This new context cannot be context 2, since in that context item 4 is not in the extension of 'is *F*'. Let's call the new context 'context 3'. Assuming that context change is conservative, context 3 is context *W*. Sentence (5) is, of course, true in context 3. So is sentence (6). So sentence (7) is true in context 3. Now – and here's where the tedium ends – we present our subject with item 4. And she knows that in context 3, sentence (7), which says that item 4 is *F*, is true. She desires to be faithful to the standards of context 3, so she explicitly endorses sentence (7). When she does so, she changes the context: the rule demands that item 5 is in the extension of 'is *F*' in any context in which item 4 has been explicitly characterized as being *F*. But this is impossible: item 5, since it is in the default antiextension of 'is *F*', is not in the

extension of 'is *F*' in any context. So it seems that Soames's solution to the sorites fails.

3. *Potential responses to the objection*

I now consider briefly two ways of responding to my objection.[5]

Response 1. Our subject has the discretion to include item 4 in the extension in some context, but she cannot ever *explicitly* exercise this discretion; she can exercise her discretion only implicitly. This is what she does by explicitly characterizing 3 as *F*.

This response is unattractive. It is not the business of a semantic theory to rule out the empirical possibility of people who are fanatically faithful to their immediately previous standards – so faithful that they feel compelled to make explicit their earlier commitments. As Soames himself says (1999, ch. 7, fn.10) whether or not a person is uncomfortable making certain explicit characterizations is a psychological, not a semantic, matter.

Response 2. Revise the view so that the range of 'everything' – but not of '*x*' and '*y*' – in the modified adjustment rule is restricted to items in the default inextension. This would involve, among other things, changing clause (iv) of the model so that adjustments involve only items in the default inextension.

To my mind the most serious problem with this suggestion is that it seems ad hoc. The plausibility of the original clause (iv) of the model seems to derive from the thought that one might contextually adjust the meaning of 'is *F*' in such a way that all things that are similar to a thing that has been explicitly characterized as *F* are counted as *F*. This sort of thought cannot underwrite the revision though. So one wonders what does.[6]

University of Kansas
Lawrence, KS 66045, USA
teresa@eagle.cc.ukans.edu

Reference

Soames, S. 1999. *Understanding Truth*. New York: Oxford University Press.

[5] Both have been floated by Soames in correspondence. I mention this only by way of thanks; the responses should not be assumed to be ones he would endorse.

[6] I should at least mention that I think there are a number of other problems that the view – whether so revised or not – faces. But discussing these problems would take me beyond my present point. I thank David Braun, Tony Genova, Nathan Salmon, Jennifer Saul, and Scott Soames. This investigation was supported by the University of Kansas General Research Fund allocation #2301791.