

April 30, 2017

April 30, 2017

Rejecting Excluded Middle

by

Scott Soames

School of Philosophy

University of Southern California

To Appear in

The Sorites Paradox

Edited by

Sergi Oms and Elia Zardini

Cambridge University Press

Rejecting Excluded Middle

1. Vagueness, Partial Definition, and Rejection of Excluded Middle

Proponents of excluded middle, $[S \text{ or } \sim S]$, typically assume that rules governing vague predicates like ‘young’ and ‘red’ are totally defined, and so determine for each object that they are true, or false, of it. Since these rules arise from ordinary uses of the predicates, this assumption raises the question of how such uses could result in distinctions, imperceptible to speakers and undiscoverable by anyone, between e.g., the last second of one’s youth and the first second at which one’s youth is merely a memory. The difficulty answering this question has led some hold that vague predicates are only partially defined, being true or false of some things and undefined for others. When P is undefined for o , neither the claim that P is true of o , nor the claim it isn’t, is sanctioned. We accept that P is (isn’t) true of o just in case we accept *that o is (isn’t) \mathcal{P} and that the claim that o is \mathcal{P} is (isn’t) true.*¹ In such cases, these claims, and the sentences expressing them, are ungrounded; they can’t be known, and even knowledge of all linguistic and nonlinguistic facts wouldn’t justify accepting them.²

Suppose that ‘is red’ is such a predicate involving the natural kind term ‘red’ standing for the property of object surfaces responsible for the fact that certain things look

¹ ‘ \mathcal{P} ’ is a schematic letter replaceable by the vague predicate that is the value of the metalinguistic variable ‘ P ’. See chapters 6 and 7 of Soames (1999) for this way of relating truth to ungroundedness.

² Pp. 364–70 of Soames (2009a.b).

similar to us and different from other things.³ On this assumption, it is a *necessary* truth that *o* is red if *o* has the physical property that *actually* explains the relevant appearances. Nevertheless ‘is red’ is partially defined because it is learned by example. Noting that nearly everyone says of things perceived to be of a given color shade, RE1, “They are red,” while saying of things perceived to be of shade RA1, “They aren’t red,” we accept the rule *Red 1*.⁴

Red

For all *o*, if *o* is RE₁, then ‘is red’ applies to *o*

For all *o*, if *o* is RA₁, then ‘is red’ does not apply to *o*

More experience leads us to adopt additional rules involving further shades until we are counted as understanding the predicate.⁵ The requirement that nearly everyone follow the same rules ensures that the language-wide rules governing it don’t determine verdicts for all possible cases.

The verdictless cases fall into several classes. Some are characterized as ‘red’ (or not red) by nearly all speakers who understand the term, even though some others withhold judgment and a few disagree. Depending on the audience, subject, and time, speakers may be more expansive in what they count as ‘red’ than they are in other contexts. Sometimes conversationalists reach implicit agreements to *count* something as ‘red’ (or not) for current purposes, while realizing that different standards may justifiably be applied in other contexts. When this occurs there is no imperative that the adopted

³ Pp. 265-66 of the reprinting of Soames (2007) in Soames (2014).

⁴ I abstract away from the complicating fact that our criteria for calling some things ‘red’ – e.g. human hair – are different from our criteria for calling other things ‘red’ – e.g. cherries.

⁵ A further factor required for mastery is agreement with one’s fellows involving comparative judgments – *x is redder than y*.

standard settle, for objects of each possible shade, whether or not ‘is red’ is true them; it is enough that it allows all conversationally relevant objects to be classified. Thus ‘is red’ will be partially defined. If ‘N’ names an object for which it is undefined, ‘N is red’ and its negation will be ungrounded; they will express propositions that can’t be objects of knowledge, and their truth or falsity won’t be determined by all linguistic and non-linguistic facts. Agents will often be indifferent about how they are classified, saying, if presented with one, *“It is and it isn’t,” “There’s no saying,”* or *“It doesn’t matter, call it what you like.”*

The fact that S and $[\sim S]$ may both be ungrounded bears on the law of the excluded middle. If a disjunction is ungrounded whenever both disjuncts are, the “law” can’t be accepted. But must it be ungrounded? It doesn’t, in general, follow from the fact that accepting each of two propositions is unjustified that accepting their disjunction is unjustified. Nor does it follow that one who fails to know each of them also fails to know their disjunction. What about determination of truth by linguistic and nonlinguistic facts? If the facts don’t determine the truth of a disjunction the disjuncts of which are ungrounded, then the disjunction will be ungrounded, and the law will fail. But that can’t be established without deciding whether $[S \text{ or } \sim S]$ is necessary, which is the point at issue.

Thus it remains to be seen whether excluded middle can be combined with partial definition. One way of doing so involves a form of supervaluationism in which one starts with a model M in which some sentences are true, some are false, and some are ungrounded. One then stipulates (i) that S is true, if S is true in all admissible bivalent extensions of M , (ii) that S is not true, if S is false in all such extensions, and (iii) that

other sentences are ungrounded. Since $[S \text{ or } \sim S]$ is true in all bivalent models, excluded middle is preserved. The difficulty is that this involves denying the apparent truism *that $[S \text{ or } R]$ can't correctly be called true unless either S or R can correctly be so called*. The story is also explanatorily tendentious. To determine whether S is true, one must first determine whether S is true in all admissible bivalent models. This presupposes a notion of *truth in a model* antecedent to supervaluationist truth and an antecedent logic used to calculate which sentences are true in which models. The idea that there is a hidden notion of truth conceptually prior to the notion of truth needed to defend the “truth” of the law of excluded middle is implausible. Further, since “classical” laws of logic are simply taken for granted, no justification for them is given.

Is there another way of combining partial definition with excluded middle that allows (1) be determinately true when (2a) and (2b) are ungrounded?⁶

1. Either N is red or N is not red.
- 2a. N is red.
- 2b. N is not red.

Let ' N^* ' be a new name designating the same object as ' N ' and let ' M ' designate a qualitative duplicate of it. Supervaluationism aside, (1), (3), and (4) should all be true or all undefined.

3. N is red or N^* is not red.
4. N is red or M is not red.

One might appeal to (5), which is reasonable because $R1$ and $R2$ are.

⁶ Here I apply an operator, 'determinately' to the ordinary (partially defined) predicate 'true' to produce the predicate *determinately true*. The operator is defined in section 4.1 of Soames (2003). Section 4.2 brings out an interesting consequence of it.

5. It is a necessary consequence of the rules of the language plus the underlying facts that substitution of 'M is red' or 'N* is red' for 'N is red' in any truth preserves truth.

R1. The status of a disjunction is entirely dependent on the status of its disjuncts.

R2. We have as much reason for taking 'N* is red' and 'M is red' to be true as we have for taking 'N is red' to be true.

But this reasoning doesn't establish the truth of (1), because R3 is as compelling as R2 when the referent of 'N' is midway between clear cases of objects of which 'is red' is true and clear cases in which it isn't.

R3 We have as much reason for taking 'N is not red' to be true as we have for taking 'N is red' to be true.

In the presence of R1, R3 leads to (6), which precludes taking (1) to be true, since it assimilates (1) to (1*), which is not evidently true, and indeed is ungrounded if 'red' is only partially defined.

6. It is a necessary consequence of the rules of the language, plus the underlying facts, that substitution of 'N is not red' for 'N is red' (or *vice versa*) in any true disjunction always preserves truth.

1* Either N is red or N is red.

In short, recognizing partially defined predicates plus ungrounded sentences and propositions precludes accepting all instances of excluded middle. Nor should one accept their negations, which are ungrounded when the instances are. These considerations lead to strong Kleene truth tables for logical connectives.

A v B		B		$\sim A$		A & B	
		T	F			T	F
A	T	T	T	T	F	T	F
	F	T	F	F	T	F	F
	U	T	U	U	U	F	U

$A \supset B$		B		
		T	F	U
A	T	T	F	U
	F	T	T	T
	U	T	U	U

$A \equiv B$		B		
		T	F	U
A	T	T	F	U
	F	F	T	F
	U	U	F	U

2. The Role of Context Sensitivity

Next we add the assumption that ‘is red’ is not only partially defined, but also context sensitive.⁷ A predicate of this sort has a *default extension* and *anti-extension*, which are the sets of things to which the language-wide rules plus nonlinguistic facts determine that it does, or doesn’t, apply. Since these sets don’t exhaust all cases, one may expand its extension or antiextension by predicating it (or its negation) of something not in either set. The rule for doing so, which is part of the meaning of the term, is that when one examines and calls such an object *o* ‘red’, and one’s hearers go along, the extension of ‘is red’ is expanded to include *o*, *plus objects discriminately redder than, or perceptually in discriminable in color from, o*. Let RE_2 be a shade that applies to precisely this class. *If an object is RE_2 , then ‘is red’ is true of it* is then a provisional rule that may implicitly be adopted.

To apply these ideas to the Sorites, we construct a sequence of *n* colored patches starting with shades that are definitely red and ending with shades that definitely aren’t (but rather are, say, orange). Adjacent patches are perceptually indistinguishable in color (when presented side by side in isolation) despite differing minutely in the physical

⁷ This view is extensively developed in Tappenden (1993), Shapiro (2006) and Soames (1999, 2002, 2003, and 2009a). For similarities and differences between Tappenden (1993) and Soames (1999) see pp. 225-26 of the latter. Context sensitivity without partial definition is advocated in Delia Graff Fara (2000).

properties responsible for color perception. Thus, they get imperceptibly less red at each step. This generates a Sorites argument.

- P1. x_1 is red.
 P2. x_1 is red \supset x_2 is red.
 \vdots
 Pn-1. x_{n-1} is red \supset x_n is red.
 C. x_n is red.

Since P1 is true, C is false, and the argument is valid, one must reject at least one premise. This is paradoxical because it seems to require saying of something, “It’s red,” while saying of its perceptually indiscernible twin, “It’s not red.” This can’t be, if, as one is inclined to think, one can’t affirm an observational predicate of something and, in the same breath, deny it of its perceptually indistinguishable twin.

The proponent of partial definition will note that one can reject a premise without denying it. If ‘is red’ is partially defined, then many of the premises are true and some are ungrounded, making it a mistake to accept them or their negations.⁸ Recognizing this blocks the derivation of C without requiring standards that divide objects of which the predicate is true from indiscernibly different objects of which it definitely isn’t. Still, one wonders whether *rejecting* (as opposed to denying) a premise, by accepting its antecedent while rejecting (as opposed to denying) its consequent, is any better. Can one reasonably say of some x_i “It’s red” while saying, in the same breath, “I reject the claim that it is” of its perceptually indiscernible twin x_{i+1} ?

⁸ See the bottom right corner of the strong Kleene truth table for ‘ \supset ’, also the table for ‘ \sim ’. When a sentence is ungrounded, the rules governing it are silent about it, and so provide no sanction for accepting it or its negation.

It is helpful to imagine a situation in which an agent A evaluates the argument dynamically, while being presented with the sequence $x_1 \dots x_n$ of colored patches one by one. Initially A sees x_1 , calls it “red,” and accepts $P1$. With x_1 in sight, x_2 is presented. Since it is perceptually indistinguishable from x_1 , A calls it “red,” accepting $P2$. Patch x_1 is then removed while x_2 remains in sight, and x_3 is displayed, which A calls “red,” accepting $P3$. At some point A applies ‘red’ to the last patch x_i in the *default extension* of ‘is red’, thereby accepting premise P_i . Although A is justified in so doing – x_i is *determinately red* after all – explicitly predicating ‘is red’ of it implicitly changes the standard to one that counts the predicate as true of the perceptually indistinguishable x_{i+1} . Recognizing the relationship between the two, A affirms ‘is red’ of x_{i+1} , tacitly changing the standard again, to one that counts ‘is red’ as true of x_{i+2} . The process may continue, with new standards adopted and more premises accepted. Eventually however, A will notice that items to which A is applying ‘is red’ are more similar to x_n , which A knows not to be red, than they are to x_1 , which is red. So, for some x_k , A will either reject or deny the characterization “It’s red”.

In so doing, A doesn’t reject or deny any proposition to which A had *explicitly* been committed. The property redness_k that A refuses to predicate of x_k differs imperceptibly from redness_{k-1} , which A had predicated of x_{k-1} , and so *implicitly* committed himself to recognizing x_k as having. Although A repudiates the standard adopted a moment earlier, there is no requirement that temporarily adopted standards not be repudiated when they are no longer useful. If A is now asked about x_{k-1} , maintaining his current standard will require treating it on a par with x_k . The stage is then set for A to move, step by step, back

toward x_1 , rejecting or denying the application of ‘red’ to items to which A previously applied it – without thereby rejecting or denying propositions previously explicitly accepted or affirmed. Every judgment A makes, from beginning to end, may be true.

Not realizing that the standards governing the use of the predicate imperceptibly change while moving through the sequence, A may find the argument paradoxical. The paradox will be resolved when A realizes that each premise is evaluated with respect to its own contextual standards, according to which it is true.⁹ Since there is no single context in which the standards governing the predicate make all the premises true, the false conclusion C is blocked.

This is easy to miss because the adjustment in standards that occurs as A moves from x_j to x_{j+1} (or conversely) is the minimum possible change that can occur at that stage. All A is asked to do is to explicitly apply ‘is red’ (or ‘isn’t red’) to something it has already been determined to be true of. Thus, the most conservative response is to apply it to the new item. Although this results in an imperceptible shift in standards, any other response would involve a bigger shift. Confusing the minimal change in standards with no change at all may leave A at a loss about which premise to reject. In fact, it doesn’t matter which premise A rejects as long as A doesn’t reject predicating ‘is red’ of an item in its default extension or ‘isn’t red’ of an item in its default anti-extension.

Some will object to the model’s requirement that there be a sharp line dividing items of which the predicate is true from perceptually indistinguishable items for which it is

⁹ A slight complication is needed to deal with the dividing line separating the last undefined item in the series leading up to the first item in the predicate’s default antiextension. The complication is explained in footnote 13 of Soames (2002).

undefined. To make the objection stick, one must distinguish the correct claim that such a line can't straightforwardly be displayed from the contentious claim that there can't be such a line. The former is a consequence of the principle CP underlying the way standards are contextually adjusted.

CP. For any two items x and y that are perceptually indistinguishable to competent observers under normal conditions, a competent agent A who affirms 'is red' of x when presented with it under such conditions is thereby implicitly committed to a standard that counts the predicate as applying to y as well. (Similarly for 'x isn't red'.)

Imagine A trying to display the line separating the last item x_j in the Sorites sequence to which 'is red' applies from the first item x_{j+1} for which it is undefined. Displaying them, A says of x_j "It's red," while rejecting a similar characterization of x_{j+1} . CP renders his remark incoherent. In predicating 'is red' of x_j , A implicitly commits himself to a standard that counts it as true of x_{j+1} , hence undermining his rejection. In applying 'is red' to x_j , A unwittingly placed the line between x_{j+1} and x_{j+2} . Nor would it help if A had said, of x_{j+1} , "the predicate is undefined for it," which would either have been false or would have moved the line again. Those who don't realize this will wrongly conclude there was no such line.

CP is also important for generalized nondynamic versions of the Sorites.

- P1. x_1 is red.
- P2. For all members x_i, x_{i+1} of the sequence $x_1 \dots x_n$, x_i is perceptually indistinguishable in color from x_{i+1} to competent observers in good light under normal conditions. Any two such x and y are the same color. So one is red iff the other is red. Hence, for each x_i of the sequence, x_{i+1} is red if x_i is red.
- C. Therefore each member of the sequence, including x_n , is red.

Since P1 is true, C is false, and the argument is valid, P2 must be rejected. Here, premises and conclusion are evaluated in a context using a single standard governing ‘is red’. In most contexts P2 will be ungrounded. Why, then, is it so seductive? In part most speakers don’t distinguish ‘is red’ not being true of o and its being undefined for o , leading them to think that they can’t maintain that an item indistinguishable from something red isn’t itself red. In addition, a seductive line of reasoning leads them to conflate CP with P2.

- i. If P2 were false then some x_i would be red even though a perceptually indistinguishable item x_{i+1} wasn’t red.
- ii. So, if I said “It’s red” of x_i and “It’s not red” of x_{i+1} , I would speak truly.
- iii. But CP doesn’t allow this; if my use ‘is red’ is true of x_i , then it is true of x_{i+1} , in which case, my use of ‘is not red’ will *not* be true of x_{i+1} .
- iv. So, given CP, P2 must be true.

Steps (i-iii) derive the nonfalsity of P2 from the truth of CP; step (iv) derives the truth of P2 from that. The latter ignores the difference between ungroundedness and falsity. But the former error is more interesting; the truth of CP doesn’t establish the nonfalsity of P2. Suppose A has seen x_{i-1} and x_{i+2} , barely discernable in color, without seeing other items $x_1 \dots x_n$. When presented with x_{i-1} and x_{i+2} , A says of x_{i-1} “It’s red” while saying of x_{i+2} “It’s not red.” Given CP, A is committed to counting ‘is red’ as true of x_i and all items preceding it, and ‘is not true’ as true of x_{i+1} and all items following it. *In such a context*, ‘is red’ is totally defined, excluded middle holds, and P2 is false, even though CP is true. Failing to see this makes the generalized argument seem paradoxical.¹⁰

CP specifies commitments undertaken by uses of a partially defined, context sensitive predicate. P2 encompasses all applications of the predicate, as used in a single

¹⁰ See ft. 11 of chapter 7 of Soames (1999).

context. The tendency to conflate claims of these kinds leads to errors beyond the Sorites. Consider the strong Kleene tables for conjunction and negation. When P is undefined for the object named by n , they determine that $[Pn]$, $[\sim Pn]$, $[Pn \ \& \ \sim Pn]$, $[\sim(Pn \ \& \ \sim Pn)]$ are ungrounded, and so rejectable. Why, then, does rejecting noncontradiction seem worse than rejecting excluded middle? Probably because *ungrounded* instances of noncontradiction are easily confused with *true* metalinguistic generalizations in ways that ungrounded instances of excluded middle aren't. When P is both partially defined and context sensitive, it is easy to confuse the ungrounded $[\sim(Pn \ \& \ \sim Pn)]$ and its ungrounded metalinguistic counterpart expressed by $[Pn \ \& \ \sim Pn]$ is *not true*, with the defensible (7).¹¹

7. No contextual standard governing P counts $[Pn \ \& \ \sim Pn]$ as true.

Similarly, confusing $[Pn \supset Pn]$ and its metalinguistic counterpart expressed by *All instances of $[Pn \supset Pn]$ are true* with the seemingly obvious truth (8) may make agents reluctant to reject instances of the former.¹²

8. If a contextual standard counts $[Pn]$ as true, it counts $[Pn]$ as true.

This explanation applies to many penumbral truths involving vague predicates. For example, the ungrounded (9a) is easily confused with the truth (9b).¹³

9a. If a man is bald, then he would be bald if he had one less hair.

¹¹ For a sentence S to be counted as true (not true) by a set of contextually adopted rules is for the claim that S is true (not true) to be a necessary consequence of the rules plus the underlying non-linguistic facts. If rules can be partial, some sentences will neither be counted as true, nor as not true. In these cases the rules are silent; for some sentences S the rules don't count S as true, and they don't count S as not true.

¹² (7) and (8) are obvious truths, provided it is obvious that P is undefined for o . If it is possible for the relation *is undefined* to fail to be defined for P and o , (7) and (8) may be ungrounded. Responses to this complication, raised in Williamson (2002), are given in Soames (2002, 2003).

¹³ See pp. 440-441 of Soames (2002).

- b. No matter what standards governing ‘is bald’ we adopt, if according to those standards *he is bald* applies to a man, then according to those same standards it would apply to him if he had one less hair.

By contrast, there is no similar truth corresponding to the law of the excluded middle that makes us reluctant to reject it.¹⁴ Consequently, the reason it seems easier to reject than other classical laws may be that rejecting it isn’t subject to the pragmatic interference we encounter with the others. Logically, the various laws have the same status. Pragmatically, they differ in what they suggest about the effects of context change.

3. The Challenge Posed by Recognizing Super-Fine Distinctions

These are ways in which some who reject excluded middle deal with challenges that come in the wake of its rejection. They maintain that in most normal contexts there are no sharp lines dividing items of which ‘is red’ counts as true from perceptibly indiscernible items of which it counts as not true. If there were, agents’ assertive commitments would be implausibly opaque, because the properties truth and falsity we use to assess them would be epistemically inaccessible in many cases. It is hard to believe that rules governing our use of language make the distribution of such important properties unknowable. The view sketched here avoids the full force of that worry. However, it does face a weakened form of it. Proponents admit there is a sharp line dividing the

¹⁴ When it is realized that rules governing predicates governing predicates may be partial, it is apparent that the metalinguistic counterpart – *Every contextual standard counts $[P \text{ or } \sim P]$ as true* – of will $[P \text{ or } \sim P]$ will be also be false. Nor will it do to suggest that the metalinguistic counterpart of $[P \text{ or } \sim P]$ is *Every contextual standard either counts P as true or doesn’t count P as true*. Although this second metalinguistic counterpart of the original is true, it is not revealing because it obliterates the partiality of the relevant predicates. For related discussion see pp. 440-4 of Soames (2002), and also Soames (2003).

default extension of ‘is red’ from perceptibly indiscernible items just outside it.¹⁵ They must explain how this line arises from the use of the predicate in the linguistic community.

The default extension and antiextension of ‘is red’ is supposed to be determined by the language-wide rule followed by all competent speakers, i.e. by all who understand it. These speakers are presumed to operate within a broad framework of agreement about items that uncontroversially count as ‘red’ or ‘not red’. Within it they are understood to have somewhat different standards for applying it. Each is disposed (i) to confidently and uniformly apply ‘is red’ (isn’t red’) to items in its default extension (antiextension), and to expect the same of others, (ii) to less confidently and uniformly apply ‘is red’ (isn’t red’) to other items, and to expect their fellow speakers to be similarly variable, and (iii) to have no consistent dispositions to apply either predicate to some further items.

If all that were true, it would make sense to posit an unknown but potentially knowable sharp line separating the default extension of the predicate from items for which it is undefined. But we can do that only if we can define what it is to be a competent speaker of the language in a way that doesn’t presuppose the picture we are trying to establish. What is it to be a competent speaker of language with partially-defined, context-sensitive predicates of the sort here described? Since, by hypothesis, the rule for ‘is red’ expresses the prescribed understanding of the predicate, it would seem that *all competent speakers* – not just a majority – should be uniformly disposed to affirm it of items in its default extension, and to affirm its negation of items in the default

¹⁵ These are items of which ‘is red’ is undefined in some contexts, and of which it is true in others only because speakers have tacitly adopted a standard that applies to them. See Soames (2003).

antiextension. We could test this, if we had a noncircular way of identifying either the competent speakers or the language they speak without presupposing the other. Given either, we could define the other. If the language were, by definition, one governed by our rule for 'is red', we could *define* a competent speaker -- as far as use of 'is red' is concerned -- to be one disposed to confidently and uniformly affirm it of all items in its default extension, and to affirm its negation of items in its default antiextension -- while being free, within certain limits, to affirm either one of other items in different contexts.

But this won't do; we aren't given the language of a group prior to knowing the linguistic dispositions of its members. Thus we must decide whose dispositions are determinative. Surely not those learning the language, who may be ignorant of common usage. Nor can we specify some percentage of a community, say 95%. There is no determining which, or how many, users of an expression are genuinely competent with it, apart from antecedent knowledge of what language we are talking about. Hence, we seem to be at an impasse. Unless we can get beyond it, the view that vague predicates like 'is red' are partially defined may face a weakened version of the same objection that is faced by the view that they are totally defined.

4. Meeting the Challenge: The Case for Micro languages

The problem seems to arise from pursuing our investigation at too high a level of abstraction. Many discussions of vagueness presuppose (i) that the investigated language is spoken by a vast but ill-defined linguistic community, (ii) that the rules governing its predicates make precise, extremely fine-grained discriminations, either between items of which they are true and those of which they are not, or between items for which they are

defined and those for which they are not, and (iii) that the contents of these rules are somehow abstracted from the totality of uses to which the predicates are put by speakers. This combination of views is hard to accept – whether or not excluded middle is rejected.

How would things look if, instead of conceiving of our (initial) object language as just indicated, we focused on the prerequisites for quick, effective, and effortless communication among members of a smaller, identifiable group speaking a micro-language the semantic properties of which are extracted from their linguistic behavior alone? For this approach to work, we would need to spell out *what it is for (non-deferential) agents to speak the same micro-language* – as far as their use of expressions like ‘is red’ and ‘isn’t red’ are concerned. Given this, we could take micro-languages to be constructions abstracted from *non-deferential* uses of expressions by groups of identifiable agents – i.e., from uses in which agents employ their internalized criteria for applying ‘is red’ or ‘isn’t red’, rather than relying on whatever standards may be employed by others. Although deferential uses of expressions may be plentiful, their contents rest on non-deferential uses. Thus the extraction of linguistic content from patterns of use should naturally focus on overlapping agreements in nondeferential dispositions to apply predicates.

Imagine a sequence of colored patches $x_1 \dots x_n$ ordered under *redder than*, starting from the reddest and ending with a patch that isn’t red (relative to some relevant color contrast). Let R_1 be the reddest shade of red, of which x_1 is an instance. R_2 , of which x_2 is an instance, is the next reddest shade. Similarly for the rest. For each adjacent pair R_i, R_{i+1} some items of which R_i is true are perceptually redder than some of which R_{i+1} is true,

some items of which both properties are true are *perceptually* indistinguishable, and no items of which R_{i+1} is true are perceptually redder than any of which R_i is true.¹⁶ We use these to specify properties implicitly associated by agents with ‘is red’. Let R^* be a property predication of which represents o as being of some shade in a particular initial segment of the sequence – *being either $R_1, R_2, \dots, \text{or } R_i$* – and predication of the negation of which represents o as being of some shade in a non-overlapping final segment of the scale – *being either $R_j, R_{j+1}, \dots, \text{or } R_n$* . R^* is true of any item of one of the initial shades, not true of any item of one of the final shades, and undefined for other items.

It is not required that R^* be totally defined, which it won’t be if the initial and final segments used to specify it don’t exhaust the sequence. It is also not required that the same property be associated by an agent with the predicate at all contexts and times. The properties associated with the predicate depend on the dispositions of agents to apply it, which may vary over time, due to the context of use and the purposes of the inquiry in which the predicate plays a role. For any agent A and time t , there will be items of which A is disposed to confidently and consistently affirm ‘it’s red’, others of which A is disposed somewhat confidently and somewhat consistently to affirm ‘it’s red’, (iii) still others of which A has no disposition to confidently and consistently affirm, either ‘it’s red’ or ‘it isn’t red’, (iv) further items of which A is disposed somewhat confidently and somewhat consistently to affirm ‘it isn’t red’, and finally items of which A is disposed to

¹⁶ As always, when speaking of items as perceptually distinguishable, or indistinguishable, I am speaking of pairwise discriminability in a situation in which they are presented together to an agent in isolation from any other items.

confidently and consistently affirm ‘it isn’t red’. Though capable of varying somewhat over time, these categories may be presumed to be reasonably stable.

We might operationalize these distinctions by tracking A’s responses -- “Yes, it’s red,” “No, it’s not red” or “I can’t say” -- to queries about items in the Sorites sequence. Imagine an n -round test, each having two parts. In part 1, A runs through an initial segment until reaching the first item at which A says “No, it’s not red,” recording any responses of “I can’t say” or “I can’t tell” that may occur along the way. In part 2, A moves backward through a final sequence and continues until reaching the first item at which A says “Yes, it’s red,” recording any responses of “I can’t say” or “I can’t tell” that may occur. Let the (partially defined) property R_A^* identified by the test be the one that is true of items for which A gave an affirmative response on every round of the test, false of items for which A always gave a negative response, and undefined for the rest. Take that be the property associated by A with ‘is red’ at t .

When the same methodology is applied to members of a group G of agents linguistically interacting with one another, R_G^* is the default property expressed by ‘is red’, which precisely determines its default extension and antiextension in a common micro-language they share – if there is such a common language. Whether or not they do depends on the nature of their overlapping dispositions to apply ‘is red’ and ‘isn’t red’.

A Necessary Condition for Speaking the Same Micro-Language

(i) *All members of G* are disposed to confidently and consistently affirm ‘is red’ (‘isn’t red’) of objects of which R_G^* is (isn’t) true (ii) They are disposed to judge items of which R_G^* is true to be “redder than” than those for which R_G^* is undefined, which the members of G are disposed to judge to be “redder than” those of which it is not true. (iii) Their dispositions impose a partial linear ordering under ‘is redder than’ of items of intermediate stages. (iv) Their disposition to affirm ‘is red’ (‘isn’t red’) of any item x_{i+1} (x_i) in the Sorites series, conditional on having

examined x_i (x_{i+1}) and affirmed ‘it’s red’ (‘it isn’t red’) of it, is very high. (Having examined and judged x_i to be ‘red’ they are temporarily disposed to judge its perceptually indiscernible twin x_{i+1} to be ‘red’ as well, and similarly for items judged ‘not red’.)¹⁷

Satisfaction of this condition provides a basis of the agreement in applying ‘is red’ needed for effective communication among (non-deferential) agents using it. But further agreement is also required. If the property R^*_A associated with ‘is red’ (‘isn’t red’) by A is true of only a fraction of the things of which the property R^*_B associated with ‘is red’ (‘isn’t red’) by B is true, while being false of many other things of which R^*_B is true, then it is reasonable to take R^*_A and R^*_B to be too different for A and B to count as speaking the same micro-language. Thus we add a further condition.

A Further Condition for Sharing a Micro-language

For all members A and B of G there is nothing of which R^*_A is true (false) and R^*_B is false (true). In short, no member of G takes some things to be paradigmatic cases of which “is not red” is true that other members of G take to be paradigmatic cases of which ‘is red’ is true.

With the addition of this condition, our conception of speaking the same micro-language requires there to be some clearly-‘red’/clearly-‘red’ agreements among agents, some clearly-‘not-red’/clearly-‘not-red’ agreements, and no clearly-‘red’/clearly-‘not-red’ disagreements – though there may be clearly-‘red’/borderline differences and clearly-‘not-red’/borderline differences.

¹⁷ Consider point (iii). Although adjacent items in the sequence are perceptually indistinguishable from one another when presented side by side, in isolation from other items, they can be expected to differ in which other items in the sequence they are pairwise perceptually discriminable from. The partial ordering will, I think, emerge when these dispositions are taken into account.

If x and y *speaking the same micro-language*, then x *speaks the same micro-language as* y and y *speaks the same micro-language as* x . But since those who *speaking the same micro-language as* x must all agree with x in certain ways and not disagree in others, and similarly for those who *speaking the same micro-language as* y , x can *speaking the same micro-language as* y , and y *speaking the same micro-language as* z , even though x *doesn't speak the same micro-language as* z . Speaking the same micro language is *sharing at least one micro-language*; typically individuals speak many micro languages differing all but imperceptibly from one another.

Think of it this way. A micro-language allows some variation among its speakers so long as it doesn't get too great. Two people "speak the same micro-language" iff there is a common micro-language they speak. (A can share a language with B, and B can share with C, even if A doesn't share with C.) A shared micro-language is an abstract object – with a phonology, syntax, and an assignment of semantic contents to expressions. Since the properties that are semantic contents of some predicates are partially defined, there are variations in how the speakers use those predicates, which means we are going to need contextually sensitive rules to determine what is asserted by speakers who address different speakers of the same micro language.

Next we define, for each individual x and time t , the notion of a *micro-language* L_{xt} *centered on* x *at* t . Agents who speak such a language L_{xt} are those who treat *at least some* items that A treats at t as paradigmatically 'red' ('not red') in the same way A does, and

who never treat something that A treats in one of those ways in the opposite way.¹⁸ The rules of this language dictate (a) that ‘is red’ applies to things that are instances of shades *all these agents* are disposed to confidently and consistently count as ‘red’, and (b) that ‘is red’ doesn’t apply to things that are instances of shades that *all the agents* are disposed to confidently and consistently count as ‘not red’. ‘Is red’ is undefined for the rest. The dispositions in terms of which these categories are determined can be given operational definitions that specify the relevant cut off points as precisely as is desired.

Next we use these ideas to illuminate the propositions asserted and beliefs expressed by uses of sentences containing ‘red’. Assume that speakers presuppose their conversational partners *speak the same micro-language as they do – i.e. that they share a micro-language with them*. Suppose this presupposition is satisfied in a conversation A is having with B. If B asks, “What color is your car?” and A answers “It’s red,” there will be four fixed points around which to construct an interpretation of A’s remark. Two are the properties $+R^*_A$ and $\sim R^*_A$, the former encompassing color-shades instances of which are A’s paradigmatically ‘red’ things and the latter encompassing color-shades instances of which are A’s paradigmatic examples of items that are ‘not red’. The second pair of interpretive fixed points are properties $+R^*_B$ and $\sim R^*_B$ that bear the same relation to B as $+R^*_A$ and $\sim R^*_A$ do to A. Either $+R^*_A$ or $+R^*_B$ will also be the property $+R^*_{AB}$ encompassing color-shades instances of which *both A and B* take to be paradigmatically

¹⁸ A stronger condition would require each agent to treat most of the shades that A takes to be paradigmatic examples of ‘red’ or ‘not red’ as A does. I leave it open whether the stronger condition is preferable to the weaker one.

‘red’. Similarly, either $\sim R^*_A$ or $\sim R^*_B$ will be the property $\sim R^*_{AB}$ encompassing color-shades instances of which *both A and B* take to be paradigmatically ‘not red’.

Though neither A nor B can identify them, these fixed points are identifiable, subject to operational precisification. Imagine orderings of subtly different shades, starting with the clearly red and progressing to the clearly not red. During a given time period, it is plausible to suppose that there is a fact of the matter about which shades an agent would characterize as “red” in each of n trials of some standard paradigm for eliciting responses; similarly for “not red.” To suppose this is to suppose that there is a property $+R^*_A$ – which, though not *identified* by A, is precisely *identifiable* -- instances of which are true of A’s paradigmatic ‘red’-examples, plus a similar *unidentified* but *identifiable* property $\sim R^*_A$ instances of which are true of A’s paradigmatic ‘not-red’-examples. The same is true of $+R^*_B$, $\sim R^*_B$, $+R^*_{AB}$ and $\sim R^*_{AB}$.

Let $(R^*_A +/\sim)$ be the *partially defined property* predication of which represents an object as bearing $+R^*_A$ and predication of the negation of which represents it as bearing $\sim R^*_A$; similarly for $(R^*_B +/\sim)$. $(R^*_{AB} +/\sim)$ is defined from the other two, and so may be identical with either of the former two properties, or distinct from both. What does A assert when, in answer to B’s question -- “*What color is your car? It isn’t red is it?*” -- A says “*Yes, it’s red*”? A and B presuppose their understandings of ‘red’ each contribute to the contents of their remarks containing it. In a give-and-take conversation like this both take what A asserts to be what B asks about, with ‘red’ contributing the same content to both. In such cases, common assertive or other speech act contents are imposed on their utterances. Thus, the proposition A asserts in answer to B’s question is not, or at least not

always, simply the proposition that predicates $(R^*_{A+/\sim})$ of the car, nor simply the one that predicates $(R^*_{B+/\sim})$.

Nor, I suspect, is it (always) the proposition that predicates $(R^*_{AB+/\sim})$ of the car. Surely, if the car is something paradigmatically red for A, A's assertion should not be rendered untrue merely because it is of a shade that would not called 'red' in *all* n of the trials in terms of which we understand B's dispositions. So long as it is one that would be classified as 'red' in most such trials, there should be no objection to accepting A's assertion. We get this result if we take *the property A to have asserted the car to have* to be an extension $E(R^*_{AB+/\sim})$ of $(R^*_{AB+/\sim})$ that is true (not true) of every item of which $(R^*_{AB+/\sim})$ is true (not true), while also being defined for every item for which $(R^*_{AB+/\sim})$ is undefined, *except those for which A and B, taken together, have no consistent dispositions to classify one way or the other*. When both A and B are positively disposed to classify items of a given shade as 'red' ('not red'), even if they would not do so on every trial, $E(R^*_{AB+/\sim})$ is true (not true) of those items. In all other cases $E(R^*_{AB+/\sim})$ is undefined. These include cases about which one or both of A and B have no definite judgment or belief, either because they tend to arbitrarily vacillate between positive and negative judgments, or because they have no such judgments to make and are indifferent about how, if at all, the items are classified. Items, if any, that A mostly, but not always, classifies as red (not red) but B mostly, but not always, classifies as not red (red) also count as undefined.¹⁹

¹⁹ When A predicates 'is red' of x in private thought, the property x is judged to have is the extension of (R^*_A+/\sim) that is true (not true) of those items that would be so classified by A on most trials. In the special case in which A is observing an item x for which the predicate is undefined in A's micro-language, and

The idea extends to linguistically communicating groups of any size. One starts with the speaker A and the group, consisting of x_1, x_2, \dots, x_n , that A addresses. Since A presupposes that there is a common micro-language that A shares with the rest of that group, the first step to determining the assertive content of A's remark is to eliminate anyone x_i who doesn't speak the same micro-language as A, as far as 'red' is concerned. Taking the remaining group, x_1, x_2, \dots, x_j , one considers the micro-language centered on A, of which they are all speakers. Interpreted as a predicate of this language, 'is red' will have an extension that is the set of things of which the property $+R^*_{A, x_1, x_2, \dots, x_j}$ is true; its anti-extension will be the set of things of which $\sim R^*_{A, x_1, x_2, \dots, x_j}$ is true. These are instances of shades that A plus each of x_1, x_2, \dots, x_j take, respectively, to be paradigmatic exemplars of 'red' and 'not red'.

The calculation determining the property $E(R^*_{A, x_1, x_2, \dots, x_j} +/\sim)$ the car is *asserted* to have can then be the same as before. If B is also a member of this larger group, then the set of items for which $E(R^*_{A, x_1, x_2, \dots, x_j} +/\sim)$ is true (not true) will either be identical with, or a proper subset of, the set of items of which $E(R^*_{AB} +/\sim)$ is true (not true). The simplest way to do the calculation is to let $E(R^*_{A, x_1, x_2, \dots, x_j} +/\sim)$ be a property that is true (not true) of all items that most of the conversational participants A, x_1, x_2, \dots, x_j are positively disposed to classify as 'red' ('not red') on most trials, even if they would not do so on every trial.

trying to decide how to (temporarily) characterize it we may take A's decision to apply 'red' to x to ensure that the property predicated of x is true of x and all items perceptually indiscriminable from x by A. Similarly for 'not red'. This special case is subject to the complication mentioned in note 9.

On this conception, the assertive contents of utterances of micro-language sentences containing the partially defined predicate ‘is red’ are context sensitive even though their semantic contents are not. Let A be a context of utterance in which A is addressing a group all of whom share a micro-language M centered on A. The lines separating the extension and anti-extension of the predicate in M from the items for which the predicate is undefined will be sharp and knowable, but almost always unknown. There need be no default extension or anti-extension because the semantic content of the predicate need not change from one context of use to another. What does change is the partially defined property contributed by the use of the predicate to the proposition asserted by the use of the sentence containing it. This may change from context to context, depending on the dispositions governing the use of the predicate by those to whom A is speaking.²⁰

If we thought that all properties had to be totally defined and all instances of excluded middle had to be true, we could revise our discussions of A and B, and of A x_1, x_2, \dots, x_j , by identifying $E(R^*_{AB} +/\sim)$ and $E(R^*_{A, x_1, x_2, \dots, x_j} +/\sim)$ with totally defined properties. In the former case, we could divide the interval between that last item of which $(R^*_{AB} +/\sim)$ is true and the first item to which it is false precisely in half, and take $E(R^*_{AB} +/\sim)_{Total}$ to be true of everything preceding the mid-point and false of everything else. The same procedure could be used for $E(R^*_{A, x_1, x_2, \dots, x_j} +/\sim)_{Total}$. In both cases, the procedure would impose sharp, imperceptible distinctions between items of which

²⁰ This may change whether or not the underlying micro-language changes, depending on the specifics of the case.

these properties are true and those of which they are false, even if no members of the relevant linguistic group have consistent dispositions to affirm or deny the properties of the items in question. In my opinion, however, this is arbitrary and unmotivated. Language users who, as a group, have no dispositions to consistently apply an observational predicate to a range of items to which they have good epistemic access, and who may not care whether or not it is applied to them are, I think, most reasonably interpreted as using the predicate to express a property that is undefined for those items.

This construction is guided by the following ideas: (i) the semantic contents of our words, and the assertive contents of uses of sentences containing them, are determined by how those words and sentences are used by agents communicating with one another; (ii) the most significant content-determining uses of a predicate like ‘is red’ are those in which sincere, non-deferential speakers affirm it, or its negation, of objects in circumstances in which they are epistemically well-placed to determine the presence or absence of the properties they judge the object to have; (iii) the semantic and assertive contents of such a predicate for a group of communicating speakers are determined by the dispositions of group members to affirm it, and its negation; (iv) effective communication among such a group, of whatever size or duration, doesn’t require complete agreement in dispositions to apply such a predicate, but it does require substantial overlapping agreements in these dispositions; and (v) because of this, the semantic and assertive contents of the words used by a non-deferential individual may change slightly depending on those with whom the individual interacts.

The construction is useful in giving us a rudimentary understanding of how the precise, knowable (though unknown) demarcations between items in the range of a vague predicate may be determined by the linguistic behavior of participants in a communicative exchange, or of a broader, but precisely identifiable, linguistic community. Although the construction can be used both by defenders of excluded middle who insist that vague predicates are totally defined, and by opponents of the law who maintain that the predicates are only partially defined, the considerations adduced here seem to favor the latter. The construction also provides a new way of thinking about the context sensitivity of (some) vague predicates. Instead of thinking ‘is red’ as having a constant Kaplanian meaning that determines (slightly) different semantic contents in different contexts, we may think of uses of it in different contexts as having (slightly) different assertive or other illocutionary contents arising from the changing mix of dispositions of speakers and their audiences, as well as, in some (but not all) cases, the slightly different semantic contents encoded by the predicate in the slightly different micro-languages employed.²¹

In applying the new scheme to the Sorites, two points should be noted. First, since partially defined predicates and properties are retained, paradox-generating premises of non-dynamic versions of the Sorites can be rejected without accepting their negations, just as before. Second, the dynamic version sketched in section 2 can

²¹ This emphasis on changes in speaker dispositions is in the spirit of Raffman (1994, 1996). Although Raffman doesn’t invoke micro-languages defined by speaker dispositions, her use of disposition change to rebut the Stanley (2003) critique of contextualist theories of vagueness can be employed formulated in terms of micro languages.

be reconstructed by trading the previous context-sensitivity of semantic content for the context sensitivity of assertive (or other illocutionary) content. Recall the featured test case in which an agent A accepts, rejects, or refuses to classify patches of color in an ordered sequence from definitely red to definitely not red in which the slight difference between an item and its successor is perceptually indiscernible to A. The judgments recorded by A's responses "Yes, that's red," ("No, that's not red"), are those indicated in footnote 19, concerning A's use of language in private thought. When A predicates 'is red' of an item x_i for which the predicate is undefined in A's micro-language, A predicates a property that is true of x_i and all items perceptually indistinguishable from it, which includes the next item in the sequence, x_{i+1} . This is all we need to get the desirable results about the dynamic Sorites mentioned in section 2.

The point is not that micro-languages fully dispose of the objection to analyses of the sort mentioned there. In order to do that, more would have to be said about the details of that construction and about the relationship between my micro-languages and the thing we call "English." If there is such a thing, it may not be a language for which any satisfying account of the Sorites can be given. The reason I haven't offered one is that I doubt that English is a well enough defined entity to support effective theorizing about the Sorites. We can, of course, talk about English in our respective micro-languages, but that talk is vague – which means that the assertive contents of our utterances containing the term 'English' are vague in the way that the assertive contents of our utterances containing the word 'red' are vague. Just as there are many different but very closely related properties that are assertive contents of various of our

uses of ‘red’, so there are many different but very closely related micro-languages that are the assertive contents of various uses of the term ‘English’.

References

- Fara, Delia Graff. 2000. “Shifting Sands: An Interest-Relative Theory of Vagueness.” *Philosophical Topics* 28: 45-81.
- Raffman, Diana. 1994. “Vagueness without Paradox.” *Philosophical Review* 103: 41-74.
- _____. 1996. “Vagueness and Context Relativity.” *Philosophical Studies* 81: 175-92.
- _____. 2005. “How to Understand Contextualism about Vagueness.” *Analysis* 65: 244-248.
- Shapiro, Stewart. 2006. *Vagueness in Context*. Oxford University Press: Oxford.
- Soames, Scott. 1999. *Understanding Truth*. Oxford University Press: New York.
- _____. 2002. “Replies.” Symposium on *Understanding Truth in Philosophy and Phenomenological Research* 65: 429-52.
- _____. 2003. “Higher-Order Vagueness for Partially Defined Predicates,” in JC Beall, ed., *Liars and Heaps*, Oxford: Clarendon Press, 128-150; reprinted in Soames 2009b.
- _____. 2007. “What are Natural Kinds?” *Philosophical Topics* 35: 329-342; reprinted in Soames 2014.
- _____. 2009a. “The Possibility of Partial Definition.” in Soames 2009b: 362-381; also in Dietz and Moruzzi, eds., *Cuts and Clouds*. Oxford University Press: New York.
- _____. 2009b. *Philosophical Essays. Vol. 2*. Princeton University Press: Princeton.
- _____. 2014. *Analytic Philosophy in America*. Princeton University Press: Princeton, 265-280.
- Stanley, Jason. 2003. “Context, Interest Relativity, and the Sorites.” *Analysis* 63: 269-280.
- Tappenden, Jamie. 1993. “The Liar and Sorites Paradoxes: Toward a Unified Treatment.” *Journal of Philosophy* 90: 551-577.
- Williamson, Timothy. 2002. “Soames on Vagueness,” *Philosophy and Phenomenological Research* 65: 422-428.