

## 7. The Context Did It

### 7.1. Survey

Many of the paradigmatically vague terms that are discussed in the vagueness literature are context sensitive. Whether the sentences *NP is tall* or *NP is rich* or *NP is bald* are true depends not just on the height, or wealth, or hairiness of the object(s) denoted by *NP*, but on the relevant standards. To borrow an example from Delia Graff, imagine Jack is both a philosopher and a Microsoft executive. When philosophers are being discussed, we can truly say *Jack is rich*. If Jack is a relatively junior executive, then when Microsoft executives are being discussed, we can truly say *Jack is not rich*. This does not mean that a contradiction, *Jack is rich and Jack is not rich*, is true. Rather, it means that the contribution the word *rich* makes to the sentence varies between the two occurrences. Sometimes we can make this variation explicit by replacing the lone word *rich* with phrases like *rich for a philosopher* or *rich for a Microsoft executive*. However, as we shall note below, perhaps not all variation in the terms can be cashed out in this way.

Given that terms have this kind of variation in content, perhaps we can use this to explain the phenomena usually associated with vagueness. What we learn from considering a Sorites series is that it is impossible for us to consider two very similar objects and judge one of them to satisfy *tall* and the other to not do so. But this does not mean that there cannot be two very similar objects such that one of them satisfies *tall* and the other does not. All it means is that the relevant pair must not be the pair we are considering. And as long as we cannot consider all the members of a Sorites series at once, it is possible that the boundary between those objects that satisfy *tall* and those that do not is somewhere that we are not looking.

We can motivate this kind of position by the following thought experiment. Imagine that we have a series of colour chips, the first of them red and the last not red, let's say it is purple, such that adjacent chips are barely discernible. (If you think that 'indiscernible' is intransitive, then you can hold that adjacent chips are indiscernible, but I don't want to take sides on that issue here.) Intuitively the boundary between *red* and *not red* is vague. One consequence of that is that we should be unwilling when presented with two adjacent chips, either simultaneously or in direct succession, to assent to *The first one is red* but not to *The second one is red*. Of course, if we do present a subject with the chips in order, she will eventually have to do just that. But we can imagine that the experiment will have the following two results. (It would be nice to know just whether the experiment does have these results. If you like, you can run the experiment for yourself, in a *very* uncontrolled setting, by downloading this file: [vaguetest.zip](#).) First, when the subject finally rejects *That's red* for one of the chips, she will be tempted to revise her immediately preceding responses to the chips. That is, once the subject decides that a chip is not red, she will be less confident in her earlier judgements that certain chips were red. Secondly, if the subject is presented with the same chips at different times, first starting with the red chips and moving to purple, and second starting with the purple chips and moving to red, she will say that more of the chips are red the first time through. The idea is that once she commits herself to a certain chip being red, this creates a rational pressure to view the next chip as red, as long as it is very similar in shade to the chip just judged to be red.

If actual experiments turn out the same way as this thought experiment, we might be tempted to conclude that subjects are prone to certain systematic kinds of perceptual errors. (Or, more accurately, to expand the range of systematic perceptual errors to which we believe subjects are prone.) The contextualist theory takes off from the idea that speakers are *not* making mistakes in these cases. Since speakers say *That is red* when presented with a particular colour patch in a red-to-purple

sequence, and say *That is not red* (or at least demur from *That is red*) when presented with the same patch in a purple-to-red sequence, and both of these things are correct, *red* must pick out a different property in the two cases. By the way, as some of the contextualists have noted, if the actual experiments do not turn out the way the thought experiments do, the contextualist theory is in a bit of trouble.

There are a few ways this basic idea can be developed. The first, and most influential, of these was developed by Hans Kamp in his 1981. The superstructure of Kamp's view is a rather semantic theory. On his view every conditional of the form *If that patch is red then so is the adjacent patch* is true, even though the quantified claim *Every patch is such that if it is red then so is the adjacent patch* is false. Since Sorites arguments are valid, a further consequence of his view is that valid arguments can have true premises and a false conclusion. These are odd results, but everyone says some odd things, so that is no reason to ignore his theory. The real reason I will say little about it here is that other theories accommodate his basic contextualist insight without such counterintuitive results, so it seems these extra costs are needless burdens.

Scott Soames (1998) takes a theory based around the three valued Kleene tables and modifies it by using these contextualist insights. Soames holds that the default meaning for a vague adjective divides objects into three categories: those that satisfy it, those that do not satisfy it, and the indeterminate, or borderline cases. Contextual effects do not shift the base cases, but they do occasionally cause us to move some of the intermediate cases into the 'true' or 'false' categories.

Building on work by Diana Raffman (1994), Delia Graff (2000) provides the most elegant of the contextualist theories. Graff's theory differs from Kamp's and Soames's in three interesting ways, and in all three cases it seems plausible that Graff is on the right side. First, Graff adopts a purely classical semantics - contextualist effects explain all the phenomena associated with vagueness, so there is no need to alter classical logic. Secondly, Graff denies that whenever we say that an object satisfies 'red' we thereby commit ourselves to holding that all similar objects satisfy 'red'. What we are committed to, apparently, is that if one object satisfies 'red' and another object is *saliently* similar, then the other object also satisfies 'red'. There is some motivation for this picture. Imagine that a school teacher has to divide her class into the A and B basketball teams. She puts the tallest students into the A team and the rest into the B team. It seems she can truly say *All and only the tall students are in the A team*. The shortest student in the A team is very close in height to the tallest student on the B team, but as long as this similarity is not salient, the teacher does not commit herself to that student satisfying 'tall'. Thirdly, Graff traces contextualist effects to explicit indexicals in the content of the vague adjectives, rather than having the contextual variation be explained along the lines of *tall for an X*. This move gives Graff a better way around the first of the objections we shall consider against the contextualist theory.

## 7.2. Problems with Context Specification

Williamson gives rather short shrift to the contextualist position in *Vagueness*. He makes just two points in order to dismiss the position. One of these is that some contextually variable terms are not vague. The example he uses, *now*, is probably not the best example. After all, unless you believe that *now* denotes a point in time, there will be some vagueness in the length of time that it denotes. And even if you do think it denotes just a point in time, there is probably some indeterminacy in figuring out just which point that is, since saying the word takes time, and it is not exactly obvious just what should count as a point in time in relativistic physics. But the underlying point here is correct: the vagueness in *now* is not due to its being context dependent. So these contextualist theories cannot just say that vagueness is context dependence. That is not a major issue. Obviously sentences are not vague merely

because we are ignorant of them. If vagueness is an epistemic phenomenon, it is a particular kind of epistemic phenomenon. Similarly, if vagueness is due to contextual variability, it is due to a particular kind of contextual variability, and not the kind displayed by *now*.

The other point Williamson makes is more substantial. Even when we specify the context, it seems that terms are still vague. For instance, *rich for an Italian poultry farmer* is a vague term, just as *rich* is vague, even though we have specified the comparison class. Indeed, there seem to be few if any ways of making the comparison class explicit that cause the vagueness to go away. If the contextualist theory is to work, it had better be because there are ways for terms like *rich* to vary that do not correspond to any ways of completing the term in question. Fortunately, there is evidence that the extension of terms like *rich* is variable in this way. Here is an example from Delia Graff that seems to illustrate the point.

Suppose I want you to hand me a certain book. If the book in question is colored a very light grayish-blue, and it's sitting amongst a bunch of other books all of which are colored a very light grayish-red, I may say, "Hand me the blue one." If, on the other hand, the book I want is sitting with a bunch of richly-colored cobalt blue books, I may say, "Hand me the gray one." I take it that it would be true to say in the first case that the book I wanted was blue, and in the second case that the book I wanted was gray. I also take it that 'gray' and 'blue' are mutually exclusive. (Graff 2000)

There's two important points to make here. First, *blue* in "Hand me the blue one" does not mean *blue for a book*. It is hard to make sense of what *blue for a book* even means. Secondly, there doesn't seem to be *any* sensible way to make the relevant comparison classes explicit in a way that reflects the contextual variation in the meanings of terms here. The best we might hope to do is to parse "Hand me the blue one" as "Hand me the one that is blue for one of those". But as Graff points out, these kinds of statements make little sense. In general, one can't say *That is blue for an X* unless the *Xs* form a kind; arbitrary groups like *those things*. So it is plausible that the contextual variation in vague terms, like *blue* or *rich* goes beyond the variation that we can generate by varying the explicit comparison classes, and hence that Williamson's objection here does not work.

### 7.3. Problems with Nouns

Graff admits that her theory is best designed to work with vague adjectives, not with vague nouns. It is hard to know just what to make of a theory of vagueness that only applies to one kind of vague term. One may have thought that systematicity was such a large virtue here that this kind of admission rules a theory out. I do not want to rest too much weight on this methodological point. Rather, I want to argue that a theory that cannot explain the vagueness in nouns cannot explain all the vagueness in adjectives that we discover.

First, let's look at some examples of vague nouns. The following two examples are both due in all essential respects to David Lewis. An old house has a carport attached to it. Now when someone points towards it and says 'that house' it is indeterminate whether they are denoting the old house, or the fusion of the old house and the carport. So the truth value of *The volume of that house is greater than  $x$  m<sup>3</sup>* might be indeterminate for some values of  $x$ . (Imagine someone says this while pointing at the house from a direction in which the carport is not visible. This removes the possibility that the indeterminacy in this case is resolved by the intention of the speaker; they will not have considered whether a carport attached in this way is part of the house.) The indeterminacy here is due to solely the indeterminacy in *that house*, all the other terms are (relatively) precise.

Proper nouns can also be vague. It is vague just which things are a part of me. Like most people, Bill Clinton occasionally loses hair. Hairs do not ‘pop’ off, they gradually fall out. At various stages of their departure, it might be indeterminate whether some hairs are part of Clinton or not. Hence it is vague just which material object is denoted by *Bill Clinton*. And this vagueness manifests in the fact that some sentences of the form *Bill Clinton’s mass is at least  $x$  kg*. For some values of  $x$ , this can be indeterminate, and this indeterminacy is, I think, traceable to the indeterminacy in the name.

The first thing to note about these cases is that it really is very hard to see how contextualist theories can account for this kind of vagueness. There is no hidden comparison class syntactically associated with *Bill Clinton*, so the vagueness cannot be generated by the fact that this class has not been generated. There is one other important disanalogy between this case and the case of vague predicates, one that does not seem to cause a problem for Soames’s version of the theory, but does raise a doubt over Graff’s and Raffman’s versions. There are clear cases of being red, but there are no clear cases of being Bill Clinton. Graff and Raffman hold the view that vague terms do not have fuzzy boundaries, they have no boundaries at all. What is surprising about vague terms, apparently, is that they classify without drawing boundaries. But that is clearly not what is going on with *Bill Clinton*. There are no objects ‘inside’ the boundary, there are just borderline cases and clear non-cases. So it is hard to see how the general story about vagueness Graff and Raffman tell, where vagueness is related to paradigmatic or typical instances, applies here, as there is no paradigmatic or typical instance of Bill Clinton.

The second thing to note is that the vagueness of these nouns threatens the completeness of the story these theorists tell about adjectives. Bill Clinton is conscious, though various objects very similar to him are not. For example, the fusion of Clinton and a hair that recently left him is not conscious. The mereological difference between Clinton and one of his loosely attached hairs is also not conscious. We can tell this by noting that when Clinton is alone in a room, there is exactly one conscious thing in the room. So just as it is vague which object is named by *Bill Clinton*, it is vague just which of the objects in the room is conscious. We know that it must be the same object that is named by *Clinton* and is conscious, but we don’t know much more than that. Since these terms are so close, if we do not have a theory that covers vague nouns like *Bill Clinton*, our theory cannot cover vague adjectives like *conscious* whose meaning is tied to the meaning of some nouns.

The main conclusion to draw from this is that if we want a theory that covers all cases of adjectival vagueness, then our theory must cover some cases of vague nouns as well. We might also note in passing that if the theory relies on there being clear positive instances of every vague adjective, then the theory must fail, because there are no clear positive instances of *being conscious*.

## 7.4. Problems with Psychologising Vagueness

Sorensen argues that Raffman’s approach to vagueness cannot work because vagueness is fundamentally not a psychological phenomenon. He also argues that Raffman’s theory conflates speaker meaning and sentence meaning. We shall look at these in order, with some attention paid to how well these kinds of criticisms affect the other contextualists.

Raffman traces the vagueness of terms like *bald* back to certain facts about how the way we judge people to be bald or not. It is not too hard to suppose that we could have the same reactions to the complex phrase *bald barber who shaves all and only those who don’t shave themselves* as we have to *bald*. We might judge that a certain hypothetical person satisfies the phrase, and that a certain different person does not, and that these two can be connected by a Sorites series. But this does not make the longer phrase vague. In fact it is perfectly precise, since it necessarily applies to no people. Since we make logical blunders, we may have the same psychological reaction to it as we have to a genuinely vague term or phrase, but that does not make it vague. The natural move here is to restrict the

contextualist theory to semantic atoms. The meaning of a complex noun phrase like Sorensen's is given by the meaning of its parts. Whether it is vague depends on how its parts, be they vague or not, interact. Whether a semantic atom is vague or precise depends on psychological factors.

This doesn't help the contextualist get off the hook. As Sorensen notes this kind of move will only work if we are prepared to trace all semantic facts back to psychological facts, and few these days will be willing to do that. (I am more willing to do it than Sorensen is, but even my internalism has its limits.) Sorensen argues that a physical ignoramus could think that we could have a Sorites series connecting an electron with a non-electron, and hence that *electron* is vague. In fact there is no such series: *electron* is a precise term. I argued in the previous section that *conscious* is vague. The argument turned on an assumption that materialism is true. If substance dualism is true, then *conscious* may well be precise. (At least it has one precise sense; it would also have a derivative sense in which it applied to bodies in virtue of those bodies being connected in the right way with souls, and in that sense it would be vague. Let us concentrate on the precise sense, however.) In that circumstance, someone who falsely believed in materialism would act as if *conscious* was vague. Still, the word *conscious* in their idiolect would be precise. Whether a word is vague or precise does not depend on our attitude towards the word.

There appears to be a further problem with tracing all vagueness back to speaker psychology: it appears to conflate speaker meaning and semantic meaning. Even if we adopt a resolutely internalist semantics, as it appears we must for these theories of vagueness to work, there is still room for such a distinction. *Her husband* denotes the husband of the person denoted by *her*, even if the speaker intends the phrase to denote some other person. It is not clear just how contextualist theories of vagueness are meant to accommodate this distinction. Assume that there are the kind of lag effects in colour perception necessary for the contextualist account to work. As a consequence of this, a speaker will sometimes classify an object as red, sometimes as not red, depending on the other objects she has seen previously. A very natural way to describe this case is that the speaker meaning of 'red' drifts even though its semantic meaning stays fixed. After all, the classic cases that Kripke describes where semantic meaning and speaker meaning come apart are cases of misperception, and lag effects seem to be perfect cases of misperception. This is not the way the contextualists wish to describe the case, but unless they do not want to recognise a distinction between semantic meaning and speaker meaning, it is hard to see what reason we could have for rejecting this account of what's happening.

## 7.5. Problems with Reports and Disagreements

There is one glaring difficulty for all of these contextualist theories. There are some rather prominent cues in language that tell us when a term has a variable content. These cues are simply not triggered for all the vague words that the contextualist claims to be contextually variable. The two cues I will focus on here concern disagreements and reports. Consider the following little dialogue.

- A:     *NP* is *F*  
 B:     *NP* is not *F*  
 C:     *A* said that *NP* is *F*

Other things being equal, we would say that *A* and *B* are disagreeing, and that *C* accurately reported what *A* said. Other things are quite prominently not equal when either *NP* or *F* is contextually variable. In the following dialogue, it is possible that *A* and *B* are not disagreeing, and that *C* is not accurately reporting what *A* said.

- A: My sister is intelligent.  
 B: My sister is not intelligent.  
 C: A said that my sister is intelligent.

We get the same phenomenon when the *NP* in question is a quantified phrase whose implicit restriction is contextually variable. For example, assume that *A* is an actor in a movie, and *B* one of the bankers for that movie. *C* asks *A* how the cast feels about the movie's progress, and *B* about how the bankers feel about its progress.

- A: Everyone is happy with its progress.  
 B: Everyone is unhappy with its progress.  
 C: A said that everyone is happy with its progress.

It seems here that what both *A* and *B* say might be true, so they are not disagreeing. If *C* says this in a context where it is possible that *A* is talking about everyone involved in the movie, then what she says is not an accurate report of what *A* said, even though she reported *C*'s words. We also get the same results with words that have an implicit comparison class. In the following case, *A* is the coach of the gym team at Katie's school, explaining why she didn't make the team, and *B* is the coach of the basketball team, also explaining why she didn't make the team, and *C* is Katie's parent complaining to *B* about the decision.

- A: Katie is tall.  
 B: Katie is not tall.  
 C: A said that Katie is tall.

Again, *A* and *B* are not disagreeing here. And it is far from clear that *C* has accurately reported *A*'s comments. So these cases are not cases where other things are equal in the above sense. Now compare this to the kind of cases that motivate Kamp, Raffman and Soames. We create duplicate sets of colour chips, with the colours moving from red to purple in barely discernable steps. *A* is shown the chips starting with a clearly red chip and moving towards purple. *B*, in a separate room, is shown the chips starting with a purple chip and moving towards red. As predicted, there are chips in the middle of the run where *A* says the chip is red, and *B* says the chip is purple. Imagine that #25 is one of these chips. So we can imagine the following things are said.

- A: Chip #25 is red.  
 (separately)  
 B: Chip #25 is not red.  
 C: A said that chip #25 is red.

It certainly seems here that *A* and *B* are disagreeing. And it seems that *C* accurately reported to *B* what *A* said. But if we have a case of contextual variation here, then this case should be just like the Katie example above. This argument is hardly conclusive; contextualists could claim that *red* is contextually variable in this way but ordinary speakers just do not recognise it in the way that they do recognise the variability in *tall*. There are limits, however, to how often this move can be made. After all, anyone can explain any recalcitrant data by this method. (The plural of *die* really is *dies*, it's just that everyone has been making a mistake about its spelling all these years.)

These problems are even more pronounced when we consider vague names. If the vagueness of *Bill Clinton* consists in it being contextually variable, then there should be contexts in which the sentences of the following form are said, and *A* and *B* are not disagreeing even though *F* means the same thing in each of their utterances.

- A*: Bill Clinton is *F*  
*B*: Bill Clinton is not *F*.

I rather doubt that satisfactory instances of this schema shall be forthcoming.

It might be noteworthy that Graff's theory does better in this respect than the other theories. Graff does not take these Sorites cases to be the paradigm cases of contextual variation. Consider again her case of the blueish-greyish book. Jack is searching through a table of books trying to find Jill's copy of *War and Peace*, which is just this colour. The left side of the table is full of bright blue books, and the right side of the table is full of grey books, the middle of the table has a few books of the same colour as *War and Peace*. *A* and *B* are in two separate booths looking down on Jack as he searches. Each of them has an obstructed view of the table; *A* can only see the grey and the blueish-greyish books, and *B* can only see the blue and the blueish-greyish books. Each of them can talk to Jack by phone, and they give him the following advice:

- A*: Jill's copy of *War and Peace* is blue.  
*B*: Jill's copy of *War and Peace* is not blue.

Jack is a little upset at this 'help' and says to *B*.

Jack: *A* said that Jill's copy of *War and Peace* is blue.

Here it is not quite so clear, I think, that *A* and *B* are disagreeing. Intuitively (at least according to my intuitions) whether the alternate colours are currently in view, as they are here, or merely have been in view in the recent past, as in the Sorites case, makes a difference. This is good news for a contextualist theory. However, this does not quite settle matters. It also seems, again defeasibly, that Jack did accurately report what *A* said. If Graff's theory is correct, it seems like this should be a correct report. So none of the contextualist theories we have considered can accommodate all the data about reports and disagreements.