

Vagueness and Practical Interest*

1. Introduction and Overview

Take the vague expression ‘tall’. Two outstanding *phenomena* of its vagueness are:

SORITES SUSCEPTIBILITY. One is inclined to accept the *soritical principle*:

(S₀) For every x and y , if x is tall and y is 1 inch shorter than x , then y is tall,

where ‘one’ encompasses at least the writers and the reader, and:

IGNORANCE OF CUT-OFFS. If y is 1 inch shorter than x , one doesn’t know that the the cut-off for being tall lies between x and y .

Clearly, a successful theory of vagueness must account for both SORITES SUSCEPTIBILITY and IGNORANCE OF CUT-OFFS.

A leading project in the contemporary philosophy of vagueness is to account for either or both of these phenomena (and possibly other ones as well) by appealing to a peculiar shiftiness exhibited by vague predicates (see e.g. Kamp, 1981; Raffman, 1994; Soames, 1999; Fara, 2000; Shapiro, 2006; Gaifman, 2007). We can usefully label all such theories ‘*contextualist theories*’. The starting point of contextualist theories is the observation that vague expressions are typically *context dependent*. Given the crucial role played in contextualist theories by the notion of context dependence, we start with a brief summary of the nature and varieties of context dependence. In general, the context dependence of an expression ε can be neutrally characterised as:

(CD) An expression ε is context dependent iff, for some contexts c_0, c_1, c_2 and c_3 , the *extension*¹ of ε as uttered in c_0 is correctly assessed in c_1 to be X while the extension of ε as uttered in c_2 is correctly assessed in c_3 not to be X .

Notice that (CD) does not imply that the extension of ε varies in virtue of the *content* expressed by ε varying from c_0 to c_2 : for all (CD) says, ε as uttered in c_2 could express the same content as ε as uttered in c_0 does and yet vary in extension. More generally, (CD)

* [Acknowledgements.]

¹ Focussing on three prominent semantic categories in philosophical logic, we assume (roughly) that the extension of a singular term is an object, the extension of a predicate a set, the extension of a sentence a truth value. Throughout, we also assume that the reader has familiarity with the broad outlines of a standard Kaplan-style semantic framework (see Kaplan, 1989).

does not imply that the extension of ε varies in virtue of features of the contexts of utterances c_0 and c_2 : for all (CD) says, c_2 could have all and only the features had by c_0 (indeed, could be numerically identical to c_0) and yet ε could vary in extension.

This neutrality of (CD) accords well with contemporary wisdom in the philosophy of language, which distinguishes at least four ways in which the extension of an expression ε can vary across contexts. Suppose that the left-hand side of (CD) holds. Then, for what we may call ‘the *standard contextualist*’, the right-hand side of (CD) holds in virtue of ε as uttered in c_0 expressing a content different from that expressed by ε as uttered in c_2 ; for what we may call ‘the *non-indexical contextualist*’, the right-hand side of (CD) holds in virtue of ε as uttered in c_0 and as uttered in c_2 expressing a single content that ε as uttered in c_0 represents as holding at a circumstance different from that in which ε as uttered in c_2 represents it as holding; for what we may call ‘the *truth relativist*’, the right-hand side of (CD) holds in virtue of ε as uttered in c_0 and as uttered in c_2 expressing a single content that is correctly assessed in c_1 to determine X as extension and correctly assessed in c_3 not to determine X as extension; finally, for what we may call ‘the *content relativist*’, the right-hand side of (CD) holds in virtue of ε as uttered in c_0 and as uttered in c_2 being correctly assessed in c_1 to express a single content different from that which they are correctly assessed in c_3 to express.

In this paper, we wish mostly to focus on a particular type of contextualist theory, according to which the context dependence that (at least partially) generates the phenomena of vagueness has its source in the variation of our *practical interests* (henceforth, for short, ‘interests’)—in other words, according to which the phenomena of the vagueness of an expression are (at least partially) due to the *interest relativity* of its correct application (we’ll call such a type of theory ‘the IR-theory’). The IR-theory has been most influentially developed and defended by Delia Graff Fara in a series of papers (Fara, 2000; Fara, 2007; Fara, 2008). Hence, in this paper we’ll largely focus on Fara’s very specific version of the IR theory—as we’ll have many occasions to appreciate, in the explanations given by contextualist theories, the devil is often in the details, and Fara is very usefully quite explicit about many of these (so, for example, we’ll see in section 6 where Fara’s version of the IR-theory exactly sits in terms of the four-way semantic divide sketched in the previous paragraph).² We suspect however that something along the lines of the details we’ll discuss is often implicit in the thoughts of those attracted by the IR-theory and by contextualist theories more generally, and so hope that our discussion will largely retain an interestingly wide scope. Thus, we’ll in effect structure our work at different levels of generality, some of these only relevant to the specifics of Fara’s version of the IR-theory while others relevant to all contextualist theories of a certain type. To anticipate our findings, we believe that one encounters fundamental difficulties at all these levels, and it is the main purpose of this paper to elaborate on what we think these difficulties are.

² A general caveat: we’ll try to be as faithful as possible to Fara’s writings, but, despite our best efforts, we might still be misrepresenting her arguments and views. Even if that were the case, we still believe that the arguments and views we do discuss have enough independent interest.

The rest of the paper is organised as follows. Section 2 discusses some arguments given by Fara in favour of the claim, arguably required by many contextualist theories, that the semantic context dependence of a certain class of adjectives cannot wholly be captured in terms of variation of comparison classes. Moving on to the theory of vagueness proper, section 3 argues against the independent plausibility of a certain salient-similarity constraint again assumed, in some form or other, by many contextualist theories. Section 4 ascends to an even higher level of generality and shows that, contrary to what seems to be presupposed by many such theories, an alleged consequence of that constraint still falls very much short of yielding satisfactory explanations of SORITES SUSCEPTIBILITY and IGNORANCE OF CUT-OFFS. Section 5 descends to a much lower level of generality, criticising Fara’s own attempt at justifying the salient-similarity constraint on the basis of a certain semantic hypothesis about the relevant class of adjectives and of considerations pertaining to our interests. Section 6 situates Fara’s version of the IR-theory in the four-way semantic divide sketched in the second last paragraph and compares its pros and cons against standard-contextualist theories along two axes: how well they score in dealing with arguments from verb-phrase ellipsis and how well they manage to preserve our pre-theoretic conception of what tallness (and many other properties expressed by vague expressions) depends on. In the face of the overwhelmingly negative findings of the previous sections for a wide class of contextualist theories, section 7 concludes by sketching two alternative approaches to SORITES SUSCEPTIBILITY and IGNORANCE OF CUT-OFFS, each favoured by (exactly) one of the authors.

2. Comparison Classes and Kinds

Fara’s version of the IR-theory is restricted to *gradable adjectives* (henceforth, ‘GAs’)³ and requires some quite substantial claims about their semantics (presupposing but significantly going beyond the main tenets of Kennedy, 1999’s theory). Before moving on in the remaining sections to see how Fara builds on these claims in her theory of vagueness, we pause in this section to expose and assess some of the more general points she makes about the semantic context dependence of GAs—in particular, her claim that such context dependence cannot wholly be captured in terms of variation of *comparison classes*.⁴ Such variation is uncontroversially crucial in explaining many cases of context dependence of GAs. For example, an utterance of ‘Jodie is tall’ made in a context c_0 where we are discussing who is to play in Jodie’s school’s basketball team is true (eight-year old Jodie is considerably tall for her age), but an utterance of the same sentence made in a context c_1 where we are discussing who is the tallest female in Scotland is false (eight-year old Jodie is not that tall). This particular variation in truth value is very naturally explained by the variation of the operative comparison class from c_0

³ Fara is of course aware that expressions belonging to other syntactic categories can also be vague (and indeed paradigmatically so, think of the noun ‘heap’). It is beyond the scope of this paper to investigate whether and how Fara’s version of the IR-theory can be extended to other syntactic categories.

⁴ In this section, we disregard world- and time-induced context dependence.

(presumably, the class of Jodie's schoolmates)⁵ to c_1 (presumably, the class of Scottish females). To show however that the context dependence of a GA cannot always be so explained in terms of variation of comparison classes, Fara, 2000, pp. 57–59 offers two different arguments.

In the first argument, which proceeds by way of *example*, we are asked to imagine the following case. We are in the West End of London and there are two different auditions taking place. One of the auditions is to find an actor suitable for playing Mikhail Gorbachev, the other is to find an actor suitable for playing Yul Brynner. Fara claims, correctly in our view, that 'bald' can be truly asserted of an actor in the context c_0 of the first audition and truly denied of the same actor in the context c_1 of the second audition. Furthermore, she claims that this is so even though the comparison class (the class of men) does not vary from c_0 to c_1 . We disagree with the latter claim. The case seems only to show that the *comparison classer* (as we may call Fara's dialectical opponent here) should assume that the relevant comparison class in the Brynner audition is a very restricted one (roughly, one containing samples from men with no hairs on their scalp to men who don't have a much larger number of hairs on their scalp than Gorbachev). That assumption seems to us independently plausible: that is the kind of comparison class that makes salient the peculiar way in which Brynner is bald.

The second argument, which proceeds in a more *theoretical* fashion, is more challenging. It goes as follows:

- (i) The comparison class of F s need not be identical to the set of things that happen to be F . It is rather, roughly, the set of things that are *typically* F ;
- (ii) The notion of a typical F requires the F s to form a *kind*;
- (iii) There are cases of context dependence of GAs where no relevant kind is available;
- (iv) Therefore, there are cases of context dependence of GAs that cannot be explained in terms of variation of comparison classes.

We may take it that the argument from (i)–(iii) to (iv) is valid. Fara motivates (i) with the following example. Imagine that, by some tragic accident, all but some very short basketball players are killed so that their average height drops drastically and Shorty is the tallest surviving one. Fara claims that, intuitively, a post-accident utterance of:

(SHORTY) Shorty is tall for a basketball player⁶

⁵ Throughout, we don't assume that the ordinary phrase 'the class of F s' (substituting an NP for ' F ') as used in theorising about comparison classes denotes the *set* of F s (or anything else that is *extensionally* individuated in terms of what things happen to be F). The ensuing discussion in this section will make clear why this assumption would be moot.

⁶ This is relevant to (i) if we assume (as we'll do) that the comparison class of F s is picked out by the PP 'for an F '.

would still be false. While we agree that one *can* understand an utterance of (SHORTY) this way, we don't think that that does anything to support (i). Let us explain. The way we see it, in the peculiar situation just described there are (at least) two salient sets: roughly, the set of pre-accident basketball players and that of post-accident ones. These two sets are everything one needs in order to explain the data, and to do so in a more satisfactory way than a friend of (i). On the one hand, there arguably is an understanding of a post-accident utterance of (SHORTY) under which it is *true*. That understanding can be explained by supposing that its truth condition is that Shorty is tall in comparison to *post-accident* basketball players. What makes that understanding not so easily accessible is the contextual salience of the set of pre-accident basketball players. And, in fact, if we change the example so as to make that set less contextually salient, any intuition of falsity subsides. Suppose, for example, that the accident happened 1,000 years ago—or, even better, that we've always been deceived into thinking that basketball players are much taller than what they actually are—with Shorty now being the tallest basketball player. Clearly, a present utterance of (SHORTY) in these two kinds of situations could be understood in such a way as to be true, and, contrary to what a friend of (i) would seem to predict,⁷ it is very doubtful that it could also be understood in such a way as to be false. On the other hand, there admittedly is also an understanding of a post-accident utterance of (SHORTY) under which, in the original situation described by Fara, it is *false*. That understanding can be explained by supposing that its truth condition is that Shorty is tall in comparison to *pre-accident* basketball players (rather than in comparison to typical basketball players, as the friend of (i) would have it). But that can in turn be easily accommodated by the independently plausible assumption that NPs carry an implicit time variable,⁸ which in the case of (SHORTY) can be contextually assigned a post-accident time (thereby generating the first reading) or a pre-accident time (thereby generating the second reading). Such an assumption does nothing to support (i).

Setting aside this particular example, it should also be noted that there are many GAs for which it should be uncontroversial that the comparison class of *F*s just is the set of things that happen to be *F*. Examples include: 'average', 'rare', 'expensive', 'poor', 'likely' etc. Hence, it should be uncontroversial that (i) is false under many substitutions for '*F*'. For all this last point shows, of course, a restricted version of the argument might still be sound. However, extra care would now be required to check that the cases witnessing (iii) do not involve one of the GAs for which (i) should uncontroversially fail. And it should always be borne in mind that the conclusion (iv) could no longer be regarded as stating a general fact about GAs.

Having said all that about (i), let's now grant it for the sake of discussion. As far as we can tell, (ii) is assumed without much argument. Given the extreme looseness of the notion of a kind, it's hard to assess (ii) in abstraction from its use in motivating other claims (like (iii)). It seems to us, however, that the notion of a kind in play should be one

⁷ At least assuming that what is typical does not change from time to time (for the first kind of situation) or from world to world (for the second kind of situation).

⁸ The point has first been emphasised by Enç, 1986, pp. ??? . A staple example are sentences like 'Every fugitive sits in jail'.

so relaxed as to be applicable, under appropriate circumstances, to just about any collection of things. Consider for example ‘big’. Take any (not all-inclusive) collection X of concrete objects. Take any (concrete) object x not in X . Consider then a rather impoverished context that has no antecedent information for the application of ‘big’. An utterance of ‘If we added x to X , we’d have another big thing’ can still be quite naturally interpreted as expressing the proposition that, if we added x to X , we’d have another thing that is big for a *typical* thing in X —i.e. a proposition where the contextually selected comparison class is intensional in the ‘typical’-sense recommended by (i), so that the utterance would be false even if x were a normal ant and X only comprised all sorts of microscopical exemplars of cars, elephants, stars etc. that happened to be much smaller than x .⁹ In this rather extreme case, the bare NP ‘typical thing in X ’ does not directly provide any *conception of typicality*, which is gleaned instead from a conception of typicality for each of the things that happen to be in X .

In less extreme cases, it is the NP itself that directly provides an at least inchoate conception of typicality. For example, Fara writes:

Comparison classes do not work just by contributing sets; for one, they need to form a kind. That is why it sounds strange to say that my computer is tall for a thing on my desk, even though it is in fact the tallest thing on my desk. Because the things on my desk don’t form a kind, we have no notion of what a *typical* height is for a thing on my desk. (Fara, 2000, p. 58)

But, despite the fact that adjoining the PP ‘for a thing on one’s desk’ to a GA results in a mild oddity, things on one’s desk can very well serve as comparison class, even in the intensional, ‘typical’-sense recommended by (i). For example, if you give a massive sequoia to a colleague as a gift and say that it is to be placed on her desk, she would certainly be within her rights in replying ‘It’s too tall’, thereby plainly expressing the proposition that it’s too tall for a *typical* thing on her desk—i.e. a proposition where the contextually selected comparison class is intensional in the ‘typical’-sense recommended by (i), so that her reply would be true even if all the things on her desk happen to be skyscrapers.

It thus seems that the notion of a kind should be appropriately relaxed, but doing this makes (iii) extremely unlikely to be true. Fara gives the example of a greyish-bluish book sometimes truly called ‘grey’, sometimes truly called ‘blue’, depending on which books surround it (whitish-bluish in the former situation, reddish-greyish in the latter). She thereby assumes that books in the surroundings do not form a kind. But again, that doesn’t seem right. Suppose for example that the books in the surroundings are books published by presses $p_0, p_1, p_2 \dots p_n$, all famous for typically producing huge books. Suppose that we’re considering whether to buy a certain standard Palgrave paperback to place in the surroundings. We would certainly be within our rights in uttering ‘That would be very small’, thereby plainly expressing the proposition that it would be very

⁹ The reader is warned that, from now until the end of this section, some fairly far-fetched possibilities will be contemplated, just to bring the point home with the utmost clarity.

small for a *typical* book in the surroundings—i.e. a proposition where the contextually selected comparison class is intensional in the ‘typical’-sense recommended by (i), so that our utterance would be true even if all the books in the surroundings happen to be miniature books exceptionally published by $p_0, p_1, p_2 \dots p_n$. This independently motivated relaxed notion of a kind allows the comparison classer to maintain that the relevant comparison class in Fara’s example just is the class of books in the surroundings. For books in the surroundings now do form a kind and, given the lack of any tight connection in the two situations described by Fara between being a book in the surroundings and having a certain colour, it is plausible to assume a conception of typicality such that, in those situations, the prevailing colour among the typical books in the surroundings just is the prevailing colour among the books in the surroundings, so that the different prevailing colours in the two situations among the books in the surroundings determine different standards¹⁰ for the application of ‘grey’ and ‘blue’.^{11,12}

We conclude that we have not seen much compelling evidence against the comparison classer. Since, as we’ll see in section 5, Fara’s version of the IR-theory arguably requires that the context dependence of GAs is not exhausted by variations in comparison classes, this raises an important issue about the independent plausibility of her approach.¹³

3. Absolute Cases and Salient Similarity

Let’s move on then to the theory of vagueness proper. Before examining which determinants of standards for the application of GAs Fara proposes to add to comparison classes, we focus in this and the next section on some crucial *constraints* that she thinks such standards must respect. For a typical GA (e.g. ‘tall’), these are (modulo relabelling):¹⁴

ABSOLUTE CASE. Certain objects must be in the extension (anti-extension) (e.g. ‘Yao Ming is tall’);

INTERNAL STRUCTURE. Extensions and anti-extensions must respect certain internal structural features (e.g. ‘Everyone at least as tall as someone tall is tall’);

¹⁰ Throughout, by ‘standards’ we’ll mean, roughly, the *thresholds* that contextually mark extensions and anti-extensions. For example, a standard for ‘tall’ is a specific height threshold.

¹¹ This possibility is of course easily missed if one assumes that the prevailing colour among the typical exemplars of a kind cannot change from world to world or from time to time. Not only is that assumption extremely implausible, but it also leads to false predictions (see fn 7).

¹² Notice that the examples of intensional comparison classes we’ve given in our discussion of (ii) and (iii) needn’t give any support to (i). For the denial of (i) is perfectly compatible with the acknowledgement that the comparison class of the *typical Fs* just is the set of typical *Fs*, and that such comparison classes are those that are contextually selected in our examples.

¹³ We don’t want to suggest however that things look particularly rosy for the comparison classer. See DeRose, 2008 for a recent attack based on some interesting examples.

¹⁴ See Fara, 2000, p. 60. The importance of the first three constraints for a theory of vagueness has first been emphasised by Fine, 1975.

EXTERNAL STRUCTURE. Extensions and anti-extensions must respect certain other extensions and anti-extensions (e.g. ‘No one is both tall and short’);

SALIENT SIMILARITY. Everything which is saliently similar to something in the extension (anti-extension) is in it.

As we’ll see, it’s SALIENT SIMILARITY that ends up doing much of the work in Fara’s version of the IR-theory. Yet, assuming classical logic,¹⁵ SALIENT SIMILARITY is in great tension with ABSOLUTE CASES (to preserve equity, we’ll see in section 5 that a crucial element of Fara’s own explanation of SALIENT SIMILARITY is also in great tension with EXTERNAL STRUCTURE and in section 6 that Fara’s own understanding of SALIENT SIMILARITY is also in great tension with INTERNAL STRUCTURE). Take a suitably fine-grained but discrete soritical series for ‘tall’. By classical logic, ABSOLUTE CASES entails there is a first case a_n to which ‘tall’ is not truly applied in any context. SALIENT SIMILARITY then recklessly entails that there is *no* context in which [‘tall’ is truly applied to the preceding case a_{n-1} and the similarity between a_{n-1} and a_n is salient].¹⁶ But what on earth is supposed to prevent the existence of such a context? It stretches credibility to maintain that the very salience of the similarity between a_{n-1} and a_n determines a shift in the cut-off between the tall and the non-tall from the loosest standard for ‘tall’ (i.e. the one that makes a_{n-1} the last positive case) to a slightly tighter one. For suppose that the loosest standard is in place, but the similarity between a_{n-1} and a_n is not (yet) salient. We say ‘ a_{n-1} is tall’ and presumably speak the truth. Suppose that we then go on to say ‘Given that a_{n-1} is tall and that a_n is very similar to it, a_n is tall too’. Following Fara, 2000, p. ???, we’ll assume throughout an intuitive understanding of salient similarity according to which *active consideration* of the similarity between two objects suffices for their similarity to be salient. Under that understanding, then, the similarity between a_{n-1} and a_n has now become salient. But how plausible is it that the last episode of speech has nevertheless managed to *tighten up* the standard for ‘tall’ in such a way that not even a_{n-1} is any more in its extension, if that speech has but re-asserted ‘ a_{n-1} is tall’? It is not that we have become more demanding in our *de facto* use of ‘tall’!

The worry generalises to a *non-classical* logical framework. For what we lack is a safety result ensuring that there is no context c where, for some suitable i , the similarity between the members of each ordered pair $\langle a_i, a_{i+1} \rangle, \langle a_{i+1}, a_{i+2} \rangle, \dots, \langle a_{n-1}, a_n \rangle$ is salient, where a_n this time is simply an absolute negative case for ‘tall’ and the determinants of standards determine that, for every $j \leq i$, a_j is a positive case for ‘tall’ in c . Keeping fixed the plausible assumption that the (few) relevant similarities can be made salient, this shows that, if SALIENT SIMILARITY holds, it will have the peculiar power of being able to trump any determination made by all the other determinants of standards. Even in a context where ‘tall’ is *de facto* used in the most relaxed way possible, with

¹⁵ As we’ll do throughout (save for a brief generalisation to be made shortly). Ditto for classical semantics (in particular, for the principle of bivalence). Hence, keeping fixed ABSOLUTE CASES, we’ll assume throughout that there are in effect cut-offs.

¹⁶ Throughout, we’ll use square brackets to disambiguate constituent structure.

everyone being most tenaciously disposed to apply it whenever ABSOLUTE CASES allows, given that the relevant similarities are salient the standard will nevertheless be quite tight!

Even if we buy into this odd metasemantic picture, related problems remain. Consider for example the series of natural numbers from 1 to 3 and a context c such that:

- (I) In c , 1 is a positive case for ‘small’ whereas 3 a negative case;
- (II) In c , it is vague whether ‘small’ applies to 2;
- (III) In c , 1 is saliently similar in size to 2 and 2 to 3.

We would have thought that there is such a c . But (I)–(III) are inconsistent if SALIENT SIMILARITY holds. This we think is bad enough a consequence of SALIENT SIMILARITY, but even if the proponent of SALIENT SIMILARITY were to stick to her guns and deny (III) (as far as we’re concerned, (I) and (II) are pretty much non-negotiable), she would be left with the embarrassment of having by her own lights an infallible method of coming to know whether ‘small’ applies or not to 2, at least assuming that, although not *saliently* similar, in c 1 is nevertheless similar in size to 2 and 2 to 3: for she knows that simply raising to salience the similarity of 2 with 1 will make ‘small’ apply to 2, while simply raising to salience the similarity of 2 with 3 will make ‘small’ not to apply to 2. SALIENT SIMILARITY would only be preserved at the expenses of IGNORANCE OF CUT-OFFS.

The last point gestures at a more general problem with SALIENT SIMILARITY. If the constraint held, it would be utterly misguided of one to entertain the hypothesis that the cut-off is where one is looking. And if one knew about the constraint, one could conclusively rule out that the cut-off is where one is looking. Given our assumption that there is in effect a cut-off (see fn 15), both conditionals strike us as having a clearly false consequent.

4. Explaining the Phenomena of Vagueness?

If it’s so problematic, why ever think that SALIENT SIMILARITY is true? Well, at least on Fara’s version of the IR-theory, it is supposed to play a crucial role in the explanation of both SORITES SUSCEPTIBILITY and IGNORANCE OF CUT-OFFS. By bivalence,¹⁷ SALIENT SIMILARITY is claimed to entail:

¹⁷ Some prominent theories of vagueness, such as the *supervaluationist* theory advocated in Fine, 1975, reject bivalence. Fara, 2000, p. ??? observes that the entailment from SALIENT SIMILARITY to TRUE INSTANCE WHEN CONSIDERED still holds if one understands SALIENT SIMILARITY as a constraint on contextually admissible, precisified, bivalent interpretations (rather than simply as a constraint on vague, non-bivalent interpretations). However, that understanding of SALIENT SIMILARITY is extremely problematic, as it would count an extremely plausible sentence like:

- (SV) A borderline case for being tall is not definitely not the cut-off for being tall

false on every contextually admissible, precisified, bivalent interpretation, and hence false (let x and y be

TRUE INSTANCE WHEN CONSIDERED. Every instance¹⁸ of a soritical principle is true when one considers it.

We say that TRUE INSTANCE WHEN CONSIDERED is only claimed to be entailed by SALIENT SIMILARITY partly because there clearly are lots of instances of soritical principles that one can consider *without* the similarity between the two objects becoming salient to one. Appealing to some far-fetched scenario, this could be argued even for instances of (S₀). But the point can plainly be made for instances of soritical principles that do not bear on their sleeves the similarity of the two objects. Suppose for example that Andy, Bill, Charlie... , Mark, Nick..., Xandra, Yetta and Zac form a decreasing soritical series for ‘tall’. Then Hero might consider ‘If x is tall and y immediately follows x in the series, y is tall’ under an assignment of Mark to ‘ x ’ and Nick to ‘ y ’ without noticing the great similarity in height between Mark and Nick. Anticipating a bit, while this point does not affect the explanation of SORITES SUSCEPTIBILITY (for one might still fully explain that phenomenon by using TRUE INSTANCE WHEN CONSIDERED suitably restricted as not to encompass subjects like Hero—after all, in her state of ignorance, Hero need not have any particular inclination to accept that conditional under that assignment), it does affect the explanation of IGNORANCE OF CUT-OFFS: it seems to us out of question that, assuming the cut-off to lie between Mark and Nick, Hero is in no better position than us to find that out.¹⁹

Conversely, another reason to doubt the claimed entailment is that it manifestly assumes ‘if’ as it occurs in a soritical principle to be some kind of *material implication* (in the sense that identity of truth values of antecedent and consequent suffices for its truth; Fara, 2000, p. ???, fn 20 seems to be aware of this glitch, but does not elaborate on it). Again anticipating a bit, while this point does not affect the explanation of IGNORANCE OF CUT-OFFS (for one might still fully explain that phenomenon by using TRUE INSTANCE WHEN CONSIDERED suitably restricted to the material-implication reading of the relevant soritical principle—after all, a cut-off is a point at which the material-implication reading of the principle fails, so that, presumably, knowing of something that it is a cut-off point requires knowing that it is a point at which the material-implication reading of the principle fails), it does affect the explanation of SORITES SUSCEPTIBILITY:

saliently similar in height and both borderline, with x taller than y : then it is true on every contextually admissible, precisified, bivalent interpretation that x is not the cut-off for being tall, and hence it is true that x is definitely not the cut-off for being tall).

¹⁸ Henceforth, we use ‘instance’ (of a quantified sentence φ) in a broad sense, so as to encompass (also) any assignment of values to the variables of φ ’s matrix. See fn 20 for a justification of this choice.

¹⁹ Puzzlingly enough, Fara might disagree. In Fara, 2000, p. 60 she gives a case where she claims that, in effect, a subject in conditions similar to Hero’s can “truly proclaim” of a cut-off that it is a cut-off. Since she explicitly contrasts that case with a case where she claims that the subject is “unable to locate” the cut-off, one might infer that Fara thinks that in her first case the subject *is able to locate* the cut-off, which in turn is usually taken to imply that the subject *knows* where the cut-off lies. If so, Fara can only uphold a watered-down version of IGNORANCE OF CUT-OFFS (one which does not encompass subjects like Hero). We find that more robust versions of IGNORANCE OF CUT-OFFS (ones which do encompass subjects like Hero) are fully warranted and find Fara’s claims about her first case very counterintuitive. Alleged examples where speakers successfully draw sharp boundaries in a soritical series actually go back to Sainsbury, 1990, pp. ????. This is not the place to delve further into such examples.

in most cases a soritical principle will be just as plausible under a reading of the conditional stronger than material implication.

Once TRUE INSTANCE WHEN CONSIDERED is in place, SORITES SUSCEPTIBILITY is supposed to be explained as follows. If one considered at a time t an instance of (S_0) , then, by TRUE INSTANCE WHEN CONSIDERED, that instance would be true at t . *Assuming* that the conditional fact about a sentence that, if one considered it at t , it would be true at t suffices for unconditionally inclining one to accept that sentence at t , one is inclined to accept any instance of (S_0) at t . *Assuming* that being inclined to accept any instance of a universally quantified sentence implies being inclined to accept the universally quantified sentence itself, one is inclined to accept (S_0) at t .

We'd first like to stress the vast implausibility of the two assumptions emphasised in the previous explanation. There are many cases where the sheer conditional fact about a sentence that, if one considered it at t it would be true at t , plainly does not promote any unconditional inclination to accept that sentence at t nor, for that matter, any conditional inclination to accept that sentence at t provided that one considers it at t . For example, supposing that Goldbach's Conjecture is true, it would indeed be the case that, if one considered it at t , it would be true at t , but that sheer fact plainly does not promote any unconditional inclination to accept Goldbach's Conjecture at t nor, for that matter, any conditional inclination to accept Goldbach's Conjecture at t provided that one considers it at t . Since presumably the unconditional inclination to accept is supposed to be secured by the conditional inclination to accept, one should at least make sure not only that every instance of (S_0) is *true* when one considers it (which is given by TRUE INSTANCE WHEN CONSIDERED), but also that one is *inclined to accept* it when one considers it. We don't find in Fara's or other contextualists' work any argument to this effect. We suspect that some *transparency* principle is implicitly being assumed, according to which we are automatically inclined to accept certain kinds of truths concerning tallness (contrary, say, to truths concerning even numbers and sums of prime numbers). We won't investigate further here to what extent some such principle might be upheld.

We prefer to emphasise that, even if some correct transparency principle allowed us to move from TRUE INSTANCE WHEN CONSIDERED to a *conditional* inclination to accept, the further move to an *unconditional* inclination to accept would still be very much unwarranted. To see this, consider for example:

(INST) For every x , if x exists, then an instance [of a universally quantified sentence] referring to x is considered by someone,

with 'if' expressing material implication. Clearly, a reflective subject will have the conditional inclination to accept any instance of (INST) when she considers it (after all, reflecting on the fact that she considers it, she will easily see that its consequent is true). But, just as clearly, there may still be lots of instances of (INST) that a reflective subject will have no unconditional inclination to accept (say, because she is not aware of the existence of the relevant object or of the relevant instance). Notice that the problem cannot be fixed by saying that the move from a conditional inclination to accept to an

unconditional inclination to accept becomes warranted under the further assumption that the subject is considering the relevant instance. For, since the move has to be warranted for each instance, this would require that the subject is considering each instance, a consequence which, together with SALIENT SIMILARITY and ABSOLUTE CASES, generates contradiction (see fn 23).

Secondly, being inclined to accept every instance of a universally quantified sentence certainly does not imply being inclined to accept the universally quantified sentence itself, for the simple reason that one might disbelieve that the instances one is inclined to accept are *all* the instances there are. It is arguable that this possibility is realised even in some cases where SORITES SUSCEPTIBILITY holds: for example, in cases where the soritical principle is known to have uncountably many instances. However, let's assume that that possibility is not realised. The main worry here is that this still falls very much short of implying that one is inclined to accept the universally quantified sentence itself. This is because of psychological versions of notorious *aggregation-failure* phenomena:²⁰ one might be inclined to accept of each participant of a fair lottery that she will lose, but one is typically not inclined to accept that everyone will lose (a psychological version of the “*lottery paradox*”, see Kyburg, 1961). Or, to give another example, one is certainly inclined to accept of each proposition about philosophy that one believes to be true that that proposition is true, but one might not be inclined to accept that every proposition about philosophy that one believes to be true is true—i.e. that one always gets it right in philosophy (a psychological version of the “*preface paradox*”, see Makinson, 1965).²¹

²⁰ If we had worked instead with the ordinary notion of an instance (roughly, the result of stripping off the initial quantifier and uniformly replacing the now free variable with a congruent interpreted expression), we would have observed also another kind of failure of that implication: in that sense of ‘instance’, one is certainly inclined to accept every instance of ‘Everything is the designation of a singular term in English’, but one is typically not inclined to accept the universally quantified sentence itself (see fn 18).

²¹ Keefe, 2007, p. 281 makes a similar point. Revising Fara, 2000, pp. ???, Fara, 2008, pp. 15–16 replies in effect by saying that the correct psychological-explanatory principle that should be assumed in the explanation is that, roughly, if one is inclined to accept the universally quantified sentence itself, that can be explained by the fact that one is inclined to accept every instance of it (rather than the principle that, if one is inclined to accept every instance of a universally quantified sentence, one is inclined to accept the universally quantified sentence itself). Setting aside broader issues in the philosophy of explanation, the new principle is even more clearly wrong than the old one: we are inclined to accept that every set of positive real numbers has a greatest lower bound, but that cannot be explained by the alleged fact that we are inclined to accept of every set of real numbers that it has a greatest lower bound—on most reasonable understandings of ‘inclined to accept’, the alleged fact does not even hold. And even if, in some sense, that fact held, the direction of explanation would certainly go the other way round. Similarly, in many cases in which one is (immodestly) inclined to accept that every proposition about philosophy that one believes to be true is true, that is explained by facts other than the completely trivial fact that one is (not immodestly) inclined to accept of each proposition about philosophy that one believes to be true that that proposition is true. Fara, 2008, p. 7 also offers the thought that the old principle might still hold defeasibly and mentions among possible defeaters knowledge that the extension of a certain predicate might have changed in the course of the reasoning and, more relevantly for our counterexamples from aggregation failure, knowledge that the conclusion has a very low probability on one’s evidence. We wholeheartedly endorse the (rather uncontroversial) defeasibility of the explanatory link (and add that insofar as the new principle is supposed to be a gloss of that, it’s a bad one). All this would however only be helpful in conjunction with an explanation—which is still lacking—of why the relevant defeaters are present in the counterexamples but not when one is assessing whether to move from (S_0) ’s instances to (S_0) itself.

Moreover, even assuming that the previous explanation works as far as it goes, we'd like to mention some cases which would seem to exhibit exactly the same kind of psychological phenomenon as that highlighted in SORITES SUSCEPTIBILITY but where SALIENT SIMILARITY and TRUE INSTANCE WHEN CONSIDERED cannot kick in. Consider for one the modification of 'tall' as 'tall by the standards of context c ', with c being a particular ordinary context. That phrase may mostly be used in philosophical English, but for all that it is a perfectly acceptable AP for competent speakers of that dialect. It is gradable. And, if y is 1 inch shorter than x , one doesn't know that the cut-off for [being tall by the standards of context c] lies between x and y . However, there is little plausibility to the idea that there is *any* context dependence in that phrase, given the usual fine-grained understanding of contexts as completely specific with regard to their agent, time and world. But SALIENT SIMILARITY and TRUE INSTANCE WHEN CONSIDERED can only hold for context-dependent expressions, on pains of reinstating the sorites paradox in the metalanguage: for example, assuming Φ to be a context-independent predicate, in a situation where, for every instance of a soritical principle for Φ , someone considers it (possibly with different subjects considering different instances at different times), SALIENT SIMILARITY would in effect act as a soritical principle for the metalinguistic predicate 'is in the extension of Φ '.²² (Indeed, as we explain in section 5, on Fara's version of the IR-theory the relevant kind of context dependence that is supposed to be at play in SORITES SUSCEPTIBILITY and IGNORANCE OF CUT-OFFS boils down to dependence on *times*. Hence, we could just as well have considered the uncontroversially ordinary AP 'tall at time t ', with t being a particular time.)

There are also cases of *quantifying* thought. Start with the soritical principle:

- (S₁) For every x , if x is small, then, for every y smaller than the second smallest natural number larger than x , y is small,

with ' x ' and ' y ' ranging over positive real numbers. Contrary to the consequent of an instance of (S₀), the consequent of an instance of (S₁) does not *refer* to any particular object—instead, it universally *quantifies* over an (uncountable) domain. Since that consequent quantifies rather than refers, it is not clear how SALIENT SIMILARITY is supposed to help with it—and in particular how SALIENT SIMILARITY could be used to derive TRUE INSTANCE WHEN CONSIDERED—for it is not clear that, when one quantifies over a set X , one needs to consider any member of X (think for example of an ordinary citizen quantifying over the set of spies). And even granting that there is in some case no obstacle in going from quantification to consideration of the objects one is quantifying over, the specifics of (S₁) present the *additional* difficulty that, in order for SALIENT SIMILARITY to be at least in the ballpark for entailing TRUE INSTANCE WHEN CONSIDERED,

²² Heck, 2003, pp. 118–120 gives voice to a widespread concern when he considers constructions that give rise to analogues of SORITES SUSCEPTIBILITY and IGNORANCE OF CUT-OFFS but which are not plausibly taken to be context dependent. While we certainly share Heck's broader worries here, none of those constructions is however a gradable AP, and hence Heck's examples do not immediately tell against Fara's version of the IR-theory, which—for better or worse—is explicitly restricted to such phrases. This explains our particular choice of example in the text.

the similarity of x with each and every real number smaller than the second smallest natural number larger than x should be salient. Setting aside the fact that it is highly controversial that normal human beings have so much as the *capacity to entertain thoughts about* many of the real numbers, it just strains credulity that *uncountably many* cases of similarity could be salient at the same time.²³

Consider next the soritical principle:

(S₂) For some F , for every x and y , if x is F and y is suitably similar to x , then y is F .

We are as much inclined to accept (S₂) as we are to accept (S₀), but it is not clear that our inclination to accept (S₂) is wholly parasitic on our inclination to accept (S₀) (or some analogous principle in which ‘tall’ is replaced by another GA).²⁴ Rather, it seems to us that one could be inclined to accept (S₂) without being aware of any witness to its truth, just as a *general* claim that properties—ways things can be—are not always sharply bounded. But such an inclination to accept could not be explained by SALIENT SIMILARITY and TRUE INSTANCE WHEN CONSIDERED.

We now turn to the explanation of IGNORANCE OF CUT-OFFS. Once TRUE INSTANCE WHEN CONSIDERED is in place, IGNORANCE OF CUT-OFFS is supposed to be explained as follows. Suppose for *reductio* that one knows at time t that [x is tall and y is not tall].²⁵ Then, by factivity of knowledge, at t x is tall and y is not tall. Moreover, one in effect knows the negation of an instance of (S₀). *Assuming* that knowing the negation of an instance at a time requires considering the instance itself at that time, one is considering the relevant instance of (S₀) at t . But then, by TRUE INSTANCE WHEN CONSIDERED, that instance is true at t . *Assuming* that the truth of the instance at t implies that at t it is not the case that [x is tall and y is not tall], we get a contradiction.

Again, we’d first like to stress the vast implausibility of the two assumptions emphasised in the previous explanation. Firstly, on no intuitive understanding of ‘considering’ does *knowing* at a time require *considering* at that time. One can know at 2.12 am on 19/04/1770 [that, [if James Cook is a man, he is not a wombat] and $2 + 2 = 4$] without considering that at 2.12 am on 19/04/1770.

²³ Considering a finite soritical series, SALIENT SIMILARITY would be straightforwardly inconsistent with ABSOLUTE CASES if all the finitely many cases of similarity could be salient at the same time. With reference to this problem, Fara, 2000, pp. ??? mentions with approval the idea that all the *finitely* many cases of similarity are too many to be all salient at the same time. But then certainly *uncountably* many cases of similarity should also count as being too many! If so, we don’t see how SALIENT SIMILARITY could be used to derive TRUE INSTANCE WHEN CONSIDERED for (S₁).

²⁴ And even if it were claimed that our inclination to accept (S₂) is wholly parasitic on our inclination to accept (S₀), this strategy would first need some refinement, since clearly inclination to accept is not generally closed under logical consequence, not even single-premise one.

²⁵ Having raised some trouble at the beginning of this section for an explicit suitable-similarity clause, henceforth we’ll oftentimes leave it out, implicitly assuming that y is *de facto* 1 inch shorter than x .

Secondly, the *semantic descent* from the instance's being true to its not being the case that [*x* is tall and *y* is not tall] is invalid given that 'tall' has an at least partially standard-contextualist semantics (as we explain in section 5, Fara would not contest this).

Moreover, points analogous to those made earlier at the end of our discussion of the explanation of SORITES SUSCEPTIBILITY can be made here (save possibly for (S₂)): there are cases which would seem to exhibit exactly the same kind of epistemic phenomenon as that highlighted in IGNORANCE OF CUT-OFFS but where SALIENT SIMILARITY and TRUE INSTANCE WHEN CONSIDERED cannot kick in. We leave the details of this to the reader.

Finally, although this is not at the centre of our interests in this paper, it should be mentioned that almost every theorist of vagueness maintains a claim much stronger than IGNORANCE OF CUT-OFFS, namely that *borderline cases in general* cannot be known by one.²⁶ It is likely that, for anyone adhering to that tenet, the *fundamental* explanation of IGNORANCE OF CUT-OFFS will be the more general explanation of the unknowability of borderline cases. For such theorists then, an explanation of IGNORANCE OF CUT-OFFS that could not be naturally extended to a more general explanation of the unknowability of borderline cases would crucially be prevented from limning the deepest features of the epistemic phenomenon surfacing in IGNORANCE OF CUT-OFFS. Needless to say, any SALIENT SIMILARITY-based explanation is, alas, one such explanation.

5. Salient Similarity and Interests

Let's take stock. Section 3 contained in effect our criticisms not only of Fara's version of a SALIENT SIMILARITY-based IR-theory, but more generally of any theory relying on SALIENT SIMILARITY. The previous section ascended to an even higher level of generality and contained in effect our criticisms not only of Fara's version of a TRUE INSTANCE WHEN CONSIDERED-based IR-theory, but more generally of any theory that tries to explain SORITES SUSCEPTIBILITY and IGNORANCE OF CUT-OFFS with TRUE INSTANCE WHEN CONSIDERED. This section descends to a much lower level of generality, investigating Fara's specific version of a SALIENT SIMILARITY-based IR-theory, and in particular the distinctive explanation she gives of SALIENT SIMILARITY on the basis of a certain semantic hypothesis about GAs and of considerations pertaining to our interests.

Fara thinks that SALIENT SIMILARITY is not a brute semantic constraint on GAs, but that it follows from their pure semantics together with some alleged facts about our interests. So let's first look at what Fara takes their pure semantics to be. Taking 'tall' as example, something like the following truth condition is proposed:

(TALL) 'x is tall' is true (at circumstances of evaluation with world *w* and time *t*) at a context with interest bearers *ii*, norm *n* and comparison class *c* iff (at *w* and *t*) *x* has significantly (for the *ii*'s interests) more height

²⁶ More precisely, that claim is stronger if one assumes (as we'll do throughout) that the cut-off is a borderline case and that there is more than one borderline case.

than n as applied to c .

So, for example, if the interest bearers are Jodie and Caitlin, the norm a norm of typicality and the comparison class the class of basketball players, ‘ x is tall’ is true (at w and t) iff (at w and t) x has significantly (for Jodie and Caitlin’s interests) more height than what is typical for a basketball player. Note that Fara understands the relativisation of truth to all the three contextual parametres in the standard-contextualist fashion: different utterances differing in these parametres will express different contents.

Let’s record a couple of worries concerning (TALL). Firstly, how can that style of analysis be *generalised* to other GAs? How can, say, ‘normal’ be analysed in terms of the ‘significantly more than’ construction? ‘ x is normal’ is true iff x has a significant amount of what? Normality? And significantly more normality than what? Than what is the norm for strange things? It strains credibility to think that we are making implicit reference to such relatively exotic comparison classes in making common-or-garden judgements of normality. Secondly, for our purposes, Fara’s analysis differs from Kennedy, 1999’s by the addition of the ‘significantly (for the ii ’s interests)’-clause. While, as we’ll see shortly, that is certainly required by Fara’s explanatory strategy, it unfortunately threatens to wreck havoc with the *exhaustivity* relations in which many GAs stand (such relations are in effect cases of EXTERNAL STRUCTURE and have first been emphasised by Fine, 1975, p. 270). So, for example, ‘red’ and ‘orange’ are arguably exhaustive over the range of the colour spectrum that goes from red to orange: if something in that range is not red, it is orange (classically equivalently, everything in that range is either red or orange). However, given this, and fixing interest bearers, norm and comparison class, analogues of (TALL) would imply that, if something in that range does not have significantly more redness than the norm, then it has significantly more orangeness than the norm. If one considers that something that does not have significantly more redness than the norm might still have more redness than the norm, and that redness and orangeness are incompatible properties, that should strike one as a problematic implication.²⁷

Having put that on the record, let’s now look at how Fara proceeds in deriving SALIENT SIMILARITY from (TALL) and its analogues together with some alleged facts about our interests. Here is her argument in a nutshell (taking ‘tall’ as example and leaving implicit interest bearers, norm and comparison class):

- (a) If two things are saliently similar (with respect to height), then they are the same for present purposes (with respect to height), i.e. for present purposes, it is fine to ignore the difference between them (with respect to

²⁷ Perhaps sensing this problem, Fara, 2000, pp. 74–75 entertains for “negative” GAs like ‘short’ a truth condition different in structure from (TALL), so that, roughly, ‘ x is short’ is true (given interest bearers ii , norm n and comparison class c) iff x *does not have* significantly (for the ii ’s interests) more height than n as applied to c . Setting aside the (pressing) questions as to whether the notion of a “negative” GA is well-defined (which we doubt) and whether the alternative truth condition is independently plausible (for ‘short’ at least it does not seem to be so), such a manoeuvre would hardly get off the ground with ‘red’ and ‘orange’, as there is no plausibility to the idea that one of them is a “positive” GA and the other one a “negative” GA and, more generally, it is very hard to see how to justify a break in their apparent symmetry by assigning to them structurally different truth conditions.

height);

- (b) If it is fine to ignore the difference between two things (with respect to height), then the costs of discriminating between them (with respect to height) outweigh the benefits;
- (c) If the costs of discriminating between two things (with respect to height) outweigh the benefits, then one is significantly taller than the norm iff the other is;
- (d) Therefore, if two things are saliently similar (with respect to height), then one is in the extension of ‘tall’ iff the other is.

Subject to a qualification to be made in fn 30, we may take it that the argument from (a)–(c) and (TALL) to (d) is valid. But, frankly, we see little justification to believe any of (a)–(c) and indeed ample justification to disbelieve them.

Take (a). It may be that we want to organise two school basketball teams for different height leagues *A* and *B*. There are 10 schoolgirls, four of which are equally very tall and four of which are equally very short. The remaining two girls, Jodie and Caitlin, are of middle height, with the one ever so slightly taller than the other one. None of the very tall girls meets the shortness criteria of the *B* league and none of the very short girls meets the tallness criteria of the *A* league; Jodie and Caitlin meet both, and no player in the *A* league can be shorter than any player in the *B* league. Given our interests in organising the two teams clearly it is not the case that it is fine to ignore the difference in height between Jodie and Caitlin: to do so would forego the only way in which we can satisfy our desire to have the two teams. And all this holds, of course, even if Jodie and Caitlin are saliently similar (with respect to height).²⁸ Such counterexamples are legion. The underlying general point is that (a) gives to the thin *contemplative* fact of salient similarity the implausible power in the *practical* domain of trumping any deeper interest we may antecedently have in not ignoring the difference between two things.

As for (b), we simply remark that it might be fine to ignore the difference between two things because the *benefits* of discriminating between them *do not outweigh* the *costs* without its being the case that the *costs* of discriminating between them *outweigh* the *benefits*. Notice that, while one could simply weaken (b) accordingly, the consequent strengthening of (c) would be even less plausible than (c).

As for (c) itself, it is hard to see why it should be true in the absence of more of an explication of what ‘significantly’ is supposed to mean. We can at least get a connection between its antecedent and consequent by interpreting the former as implying that the difference in height between two things is not significant. But it would be a fallacy in the *logic of significance* to think that that in turn implies the consequent of (c). In general, that the difference between *x* and *y* is not significant and that the difference between *x* and *z* is significant does not entail that the difference between *y* and *z* is significant.

²⁸ Henceforth, we’ll omit this qualification.

Finally, it is interesting to ask oneself why, if the (a)–(d) argument is any good, an analogous argument should not go through when *mere* similarity is substituted for *salient* similarity in (a), thus yielding:

- (a') If two things are similar (with respect to height), then they are the same for present purposes (with respect to height), i.e. for present purposes, it is fine to ignore the difference between them (with respect to height).

Such an analogous argument would of course be unacceptable, as it would reinstate the sorites paradox. Fara seems to want to block this new argument at (a'): the fact that two things are similar is supposed not to imply that, for present purposes, it is fine to ignore the difference between them. However, we don't see why mere salience should make a difference as to whether, for present purposes, it is fine to ignore the difference between them: as we've already mentioned in our discussion of (a), salience is too thin a contemplative fact to have such wide-ranging practical effects. In support of her line of thought here, Fara offers the further claim that, if two things are not "live options", then there is no cost in discriminating between them. Again, we don't see why the thin contemplative fact of salience should have the substantial practical effect of making two things "live options": at least in the ordinary sense of 'live option', not every option that is actively considered is a live option (and not every live option is actively considered). We also don't see why not being a "live option" should imply that there are no costs in making the relevant discrimination—if anything, we would have thought that it implies that there are rather no benefits (and some costs) in making the relevant discrimination! (Together with premise (c), that would of course reinstate the sorites paradox.)

More generally, it is unclear to us that it is not the case that, for present purposes, it is fine to ignore the difference between any two similar things and that, for any two such things, the costs of discriminating between them outweigh the benefits. Both these claims strike us as correct in a large variety of cases. And, given the manifest *non-transitivity* of the relations they employ, neither claim reinstates by itself the sorites paradox. Both do, however, in conjunction with (b), (c) and (TALL). Being wedded to the latter three, Fara tries to parry this attack by observing that the costs of discriminating somewhere do not outweigh the benefits of doing so. We wholeheartedly agree and would go even further by saying that the benefits of discriminating somewhere outweigh the costs of doing so, but such a platitude does little to ward off the attack. For one need not make the discrimination between similar objects—one may rather do it between non-similar objects (again, the non-transitivity of non-discrimination allows that, in general, one may discriminate between x and y without discriminating between any objects strictly between x and y).²⁹

²⁹ The non-discrimination in question is *practical* rather than *epistemic*, but it is still extremely plausibly taken to be non-transitive. The point is easily missed by taking non-discrimination between x and y to consist in x 's being treated in the *same* way as y and then fanatically assuming that the latter relation involves a transitive-identity relation. It clearly doesn't, as x 's being treated in the same way as y tolerates minute and insignificant differences in their treatment, just as x 's walking in the same way as y tolerates minute and insignificant differences in their gait. One could of course stipulate to be using a notion of non-

6. The IR-Theory vs Standard-Contextualist Theories

We finally turn to a comparison of Fara's version of the IR-theory with other contextualist theories which try to explain the phenomena of vagueness by variously appealing to context dependence understood along *standard-contextualist* lines (see e.g. ???). Although, as has emerged in the previous section, Fara's analysis too postulates several elements of standard-contextualist context dependence in the semantics of GAs (interest bearers, norm, comparison class), the previous section should also have made clear that it is not these elements that are exploited in her explanation of SALIENT SIMILARITY (which, to recall, is in turn supposed to yield TRUE INSTANCE WHEN CONSIDERED which is in turn supposed to yield SORITES SUSCEPTIBILITY and IGNORANCE OF CUT-OFFS). For what her explanation exploits is rather the fact that salient similarity to the *iis* between x and y implies that the *ii*'s interests are in a certain way, a way that ensures that x is in the extension of a GA iff y is (again, recall the (a)–(d) argument of the previous section).³⁰ Thus, the relevant changes that are supposed to ensure SALIENT SIMILARITY are changes neither of the interest bearers, nor of the norm, nor of the comparison class, but *changes in what the interest bearers' interests are*. These changes are due to what is saliently similar to the interest bearers and concern what pairs of objects is in their interest to discriminate. Fara further assumes a *temporalist* semantics, so that the relevant changes do not induce any change in the content expressed by a GA; on an *eternalist* semantics, the only change in content would be that of the time involved in the content. For simplicity, we'll also presuppose a temporalist semantics, but

discrimination that, contrary to the ordinary one, requires perfect match, but then (b) would be blatantly false.

³⁰ Actually, this points to some further wrinkles in the explanation of SALIENT SIMILARITY. In the previous section, we induced the reader not to think too hard about interest bearers by saying that these had been left implicit in the (a)–(d) argument. However, the step from (a)–(c) and (TALL) to (d) is only valid if the interest bearers that are considering x and y and are referred to throughout (a)–(c) are the *same* as those that are semantically relevant and are referred to in the operative instance of (TALL). In Fara's and other contextualists' work, this issue is constantly obfuscated by an indiscriminating use of the first-person plural pronoun. We would like to see some justification however for the required assumption that, in using a GA, the value of the interest-bearers parametre necessarily *contains the speaker*, especially given that similar assumptions about "speaker inclusion" are false (the domain to which 'everyone' is contextually restricted needn't contain the speaker, the body of knowledge relevant to 'might' needn't include the speaker's, 'I' itself needn't refer to the speaker). Moreover, even if the interest bearers were necessarily to contain the speaker, once one considers that they presumably may include also other people, the question must arise as to why, if the similarity of x and y is salient to the speaker, the discrimination between them could not be made by some *other* of the interest bearers to whom such similarity is not salient, in such a way that, after all, the costs of discriminating between them would not outweigh the benefits. To admit as a group of interest bearers at a time only groups of people in which everyone at that time is considering the same objects seems to us hopelessly *ad hoc*. With respect to this specific set of problems, Fara's version of the IR-theory seems to us to work smoothly only if one assumes a relatively *solipsistic* view of the contextual interpretation of GAs, such that the value of the interest-bearers parametre fully coincides with the speaker (or at best by the interests of those who are considering the same objects as the speaker). This seems to us to be in tension with the insights in the literature about the *communal* contextual interpretation of GAs (as found e.g. in Lewis, 1979). We're indebted here to conversations with Laura Delgado.

emphasise that the choice is irrelevant for the substance of our discussion in this section (see fn 31).

It is this relative *invariance in content* that helps Fara's version of the IR-theory to avoid some of the arguments against standard-contextualist theories from the content-invariant interpretation of context-dependent expressions under *verb-phrase ellipsis* (see Stanley, 2003). Assuming (as we do throughout) the content-invariant interpretation of context-dependent expressions under verb-phrase ellipsis, Stanley argued that standard-contextualist theories cannot explain SORITES SUSCEPTIBILITY and IGNORANCE OF CUT-OFFS for the following sorites:

(VPES) This man is tall; and this₁ is too; and this₂ is too... ; and this_n is too,

where 'this_i' [$1 \leq i \leq n$] refers to a man just slightly shorter than the previous one and 'this_n' refers to a clearly short man. This is so because the content of the elided verb phrase is invariant throughout (VPES), and so no variation in content is available, contrary to what would be required by a standard-contextualist theory. By contrast, Fara's version of the IR-theory (with (TALL) understood temporally) has no problem with (VPES), since, even though the *content* of the elided verb phrase cannot change throughout (VPES), our *interests* can, and with them the *extension* of the elided verb phrase.³¹

This is fine as far as it goes, but we don't think it goes very far. As Fara's version of the IR-theory stands at this point, it cannot explain SORITES SUSCEPTIBILITY and IGNORANCE OF CUT-OFFS for the following *modal* sorites:

(MS₀) This man is tall and interests have not such-and-such maximally specific property; but if they had, the man would still be even if he were just slightly shorter than that₁; and if they had, the man would still be even if he were just slightly shorter than that₂... ; and if they had, the man would still be even if he were just slightly shorter than that_n,

where 'that_i' [$1 \leq i \leq n$] refers to the height exemplified by the man in the previous counterfactual circumstance and 'that_n' refers to a clearly short man. This is so because the ellipsis on the verb phrase 'have such-and-such maximally specific property' ensures that in all the various counterfactual circumstances entertained throughout (MS₀) the interests cannot change, while the ellipsis on the verb phrase 'is tall' ensures that the values of the standard-contextualist parameters of 'tall' cannot change throughout (MS₀). Together, these two invariances force not only the content expressed by 'tall' throughout

³¹ An eternalist analogue of Fara's version of the IR-theory would instead require that the content of the elided verb phrase changes with respect to the time it involves. There is independent reason to think that that is possible ('Andy is in the room'; after observing Bill coming in, 'And Bill is too'; after observing Charlie coming in, 'And Charlie is too'... ; after observing Zac coming in, 'And Zac is too'). For those who think that in such cases the time is really contributed by the verb, there is also independent reason to think that the same kind of change is possible even in the case of *bare-argument ellipsis* ('Andy is in the room'; after observing Bill coming in, 'And Bill too'; after observing Charlie coming in, 'And Charlie too'... ; after observing Zac coming in, 'And Zac too').

(MS₀) to remain constant, but also its extension relative to the various counterfactual circumstances entertained throughout (MS₀), contrary to what would be required by Fara’s version of the IR-theory.

Indeed, the trick of (MS₀) is really that of forcing, in the vanilla way afforded by counterfactuals, the evaluation of ‘tall’ to be made relative to a circumstance where the interests are not guaranteed to be the actual ones of the speaker (see fn 30 for less vanilla ways of achieving similar results). It is no surprise then that a similar point could be made without exploiting verb-phrase ellipsis to keep the interests constant, as in the following *modal sorites*:

(MS₁) This man is tall; and he would still be even if he were just slightly shorter than that₁ and interests were just slightly different from the way they now actually are;³² and he would still be even if it were just slightly shorter than that₂ and interests were just slightly different from those_{1...} ; and he would still be even if he were just slightly shorter than that_n and interests were just slightly different from those_{n-1},

with the same notational conventions as before plus the conventions that ‘those_i’ [$1 \leq i \leq n - 1$] refers to the interests obtaining in the previous counterfactual circumstance and ‘those_{n-1}’ refers to clearly possible interests. Although interests are now allowed to change from one counterfactual circumstance to another, their change is, as it were, fixed and cannot be further influenced by salient-similarity facts, contrary to what would be required by Fara’s version of the IR-theory (for good measure, we still use verb-phrase ellipsis to ensure that such facts cannot even affect the values of the standard-contextualist parameters of ‘tall’).³³

Thus, we don’t think that Fara’s version of the IR-theory has a *substantial advantage* over standard-contextualist theories as far as matters of verb-phrase-ellipsis arguments are concerned. Indeed, similar considerations of modal embeddings lead us to observe a *clear disadvantage* of Fara’s version of the IR-theory over standard-contextualist theories. For suppose that the boundary between the positive and the negative cases of ‘tall’ (as uttered by speaker *s* at world *w* at time *t*) falls between *x* and *y*. Then, on Fara’s version of the IR-theory, the conditionals ‘If I considered (consider) the similarity in height between *x* and *y*, either *x* would (will) cease to be tall or *y* would (will) become tall’ are true (as uttered by *s* at *w* at *t*; we’ll assume here the solipsistic view mentioned at the end of fn 30). We find those conditionals extremely unpalatable and deem them to embody a severe misconception of *what tallness depends on*: not only as depending on someone’s considering things, but also as depending on *s* rather than anyone else considering things! Standard-contextualist theories, even if SALIENT

³² In a particular way (which in the text we don’t bother to specify). Assuming that interests are fully determined by which pairs of neighbouring objects one is considering (and that one is considering exactly one such pair at a time), we can think that successive pairs are considered in successive counterfactual circumstances.

³³ As one might expect, analogous *temporal sorites* can be constructed. We leave the details of this to the interested reader.

SIMILARITY-based, do remarkably better on this score, since, at least on the usual way of developing them, they allow all suchlike conditionals to be false. What they do require to be true (as uttered by s at w at t) are rather more theoretical conditionals such as ‘If I considered (consider) the similarity in height between x and y , either x would (will) not fall under the extension of ‘tall’ as uttered by me then or y would (will) fall under the extension of ‘tall’ as uttered by me then’. The argument extends to many other conditional and non-conditional sentences that are extremely unpalatable, but true on Fara’s version of the IR-theory, like ‘It is possible that something becomes tall without changing height’.

Indeed, even INTERNAL STRUCTURE is now in jeopardy. Suppose that the cut-off for being tall in a certain soritical series lies between Andy (who is 6 feet and 2 inches tall) and Bill (who is 6 feet and 1 inches tall), and that one is considering another series where the similarity between Charlie (who is also 6 feet and 2 inches tall) and Dave (who is also 6 feet and 2 inches tall) is salient. Then, on Fara’s version of the IR-theory, one could truly utter ‘Either Charlie is not tall and is at least as tall as someone who is tall or Dave is tall and someone [who is at least as tall as he is] is not tall’, which is straightforwardly inconsistent with our (and Fara, 2000, p. ???’s) paradigmatic example of INTERNAL STRUCTURE. Fara, just as well as a standard contextualist relying on SALIENT SIMILARITY, might rejoin that SALIENT SIMILARITY should be charitably understood as pertaining to any pair of objects that agrees in the relevant supervenience base with the pair whose similarity is salient. We don’t see a lot of independent plausibility for such a move, especially if SALIENT SIMILARITY is supposed to be explained in something like the praxis-oriented way attempted by Fara, but let that pass. Suppose that the cut-off lay between Andy and Bill at time t_0 and the similarity between Charlie and Dave becomes salient at a later time t_1 . Then, on Fara’s version of the IR-theory but not on a standard-contextualist theory, one could truly utter at t_1 ‘Either Charlie is not tall and is at least as tall as someone [who (at t_0) was tall] was (at t_0) or [Dave is tall and someone [who (at t_0) was at least as tall as he is] was not tall (at t_0)]’, which is straightforwardly inconsistent with the platitudes that a man who is at least as tall as [a man who was tall was] is tall (first disjunct) and that a man who was at least as tall as [a man who is tall is] was tall (second disjunct), platitudes which we take also to belong to INTERNAL STRUCTURE.

It is sometimes suggested that the last series of counterintuitive results can be avoided by the IR-theory if the semantics “*rigidifies*” the speaker’s actual and present interests (see e.g. Stanley, 2003, pp. ???), as in:

(TALL’) ‘ x is tall’ is true (at circumstances of evaluation with world w and time t) in a context with interest bearers ii , norm n , comparison class c , world w' and time t' iff (at w and t) x has significantly (for the ii ’s interests as they are at w' and t') more height than n as applied to c .

We don’t think however that, in this dialectical context, rigidification would be a wise move on behalf of the IR-theory. Firstly, as far as the conception on what tallness depends on is concerned, even a rigidified version of the IR-theory would be stuck e.g.

with the truth of ‘If I’m considering the similarity in height between x and y , x is tall iff y is, even though as a matter of fact x is tall and y is not’ (as uttered by s at w at t). The initial problem with INTERNAL STRUCTURE would also remain. Secondly, as far as arguments from verb-phrase ellipsis are concerned, a rigidified version of the IR-theory would have the same problem with (VPES) as standard-contextualist theories have, since, although our interests can change throughout (VPES), the invariant content of the elided verb phrase would always refer back to the interests we had at the time of utterance of the first sentence of (VPES).

We conclude that, as far as the issues discussed in this section are concerned, neither Fara’s nor a rigidified version of the IR-theory offers substantial advantages over standard-contextualist theories and that, on the contrary, they suffer from some significant disadvantages in that they license claims that do violence to our ordinary conception of what tallness depends on.

7. Conclusion

Beyond issues of detail pertaining to the specifics of Fara’s version of the IR-theory, there are two main points that we hope to have impressed on the reader: the implausibility of the semantic and psychological assumptions that one needs to make in order to get a version of the IR-theory with some hope to engage with the phenomena of vagueness (such as e.g. a version equipped with SALIENT SIMILARITY) and the great gap that would actually still remain between such version and the phenomena of vagueness. Although this is a topic for another paper, we believe that analogues of these two points hold for a wide variety of contextualist theories. If so, other, in certain respects more radical approaches to the phenomena of vagueness begin to look more appealing. One such approach disputes the status of sacrosanct data that we have accorded so far to SORITES SUSCEPTIBILITY and IGNORANCE OF CUT-OFFS: it argues that it is exactly focus on the context dependence of vague expressions that allows us to see how one can quite competently draw in context a perfectly sharp boundary between positive and negative cases (see Sweeney, 2010). At the opposite end of the spectrum, another such approach takes those two phenomena much more seriously than most contextualist theories do: by revising classical logic, it declares true soritical principles such as (S_0) , from which it is then able to derive explanations of SORITES SUSCEPTIBILITY and IGNORANCE OF CUT-OFFS that exhibit a shocking simplicity in comparison with the convolutions and epicycles with which contextualist theories are typically saddled (see Zardini, 2008; 2010).

That languages spoken by human beings contain vague expressions at least partly because that allows them to pursue certain practical interests which could not otherwise be (easily) pursued is a non-obvious but appealing hypothesis to which we’re very sympathetic. But it is not the IR-theory. That theory rather holds that what underlies certain phenomena of vagueness is the variability of those interests. That is what we have argued against in this paper. On the one hand, no sane man’s practical interests are as volatile as they’d have to be in order to generate the required variability; on the other

hand, that great riddle surfacing in SORITES SUSCEPTIBILITY and IGNORANCE OF CUT-OFFS reaches much deeper into our thoughts about reality than anything that any such variability could ground.

References

- Beall, JC (ed.) [2003], *Liars and Heaps*, Oxford UP, Oxford.
- DeRose, K. 'Gradable Adjectives: A Defense of Pluralism', *Australasian Journal of Philosophy*
- Enç, M. [1986], 'Towards a Referential Analysis of Temporal Expressions', *Linguistics and Philosophy*, pp.405-426
- Graff-Fara, D. [2000], 'Shifting Sands: An Interest-Relative Theory of Vagueness', *Philosophical Topics* 28, pp. 45–81.
- Fara, D. 2007 'Profiling Interest Relativity'
- Fara, D. 2008 'Generalizing from the Instances within "Contextualist" Accounts of Vagueness' ms.
- Fine, K. [1975], 'Vagueness, Truth and Logic', *Synthese* 30, pp. 265–300.
- Gaifman 2007
- Graff-Fara, D. [2000], 'Shifting Sands: An Interest-Relative Theory of Vagueness', *Philosophical Topics* 28, pp. 45–81.
- Heck, R. [2003], 'Semantic Accounts of Vagueness', in Beall [2003], pp. 106–27.
- Kamp, H. [1981], 'The Paradox of the Heap', in Mönnich [1981], pp. 225–77.
- Kaplan, D. 1989, 'Demonstratives
- Keefe, R., Smith, P. (eds.) [1997], *Vagueness: A Reader*, MIT Press, Cambridge MA.
- Keefe, R [2007] 'Vagueness without Context Change', *Mind* 116 (462), pp. 275-292.
- Kennedy, C. [1999] *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*, Garland Press, New York.
- Kyburg, H. [1961] *Probability and the Logic of Rational Belief*, Middletown, CT, Wesleyan University Press.
- Lewis, 1979, 'Scorekeeping in a Language Game'
- Mönnich, U. (ed.) [1981], *Aspects of Philosophical Logic*, Reidel, Dordrecht.
- Makinson, D. [1965] 'Paradox of the Preface', *Analysis* 25, pp. 205-207
- Raffman 1994
- Shapiro 2006
- Soames 1999
- Stanley, J. [2003], 'Context, interest relativity and the sorites', *Analysis* 63, pp. 269–80.
- Sweeney, 2010
- Zardini 2008
- Zardini 2010