

## Profiling Interest Relativity

DELIA GRAFF FARA

Princeton University

Final Draft, July 2008

(Forthcoming in *Analysis*)

One way to object to a theory is to argue that it has bad consequences. One way to rebut such an objection is to show that the theory does not really have the consequences being attributed to it. An equally good way to rebut such an objection is to show that the consequences are not really that bad. Here I rebut a two-part objection to my interest-relative theory of vagueness in both of these ways.<sup>1</sup> The objection, as developed by Jason Stanley (2003, 2005), concerns the modal and epistemological profiles of interest-relative propositions. Stanley's claims about the epistemological profile of interest-relative propositions are incorrect. His claims about the modal profile of interest-relative propositions are correct, but not worrisome.

### 1. The View under Attack

I hold that vague predicates express interest-relative properties. A property is *interest relative* if whether or not a thing has it depends, at least in part, on whether some things' interests are furthered or satisfied. The extension of an interest-relative property might change because (i) interests themselves change, or because (ii) what satisfies an unchanging interest changes as a result of external factors: (i) I could change from preferring a tall cherry tree to preferring a short one; or (ii) the height of a given tree could change, so that I now prefer it, whereas I did not before. On my view, the vague predicate 'is tall for a cherry tree' is to be analysed as *has significantly more height than is the norm for a cherry tree*, and the predicate 'is tall' is to be analysed as *has significantly more height than is the norm for C*, where *C* is some contextually-given comparison class. These are interest-relative properties since *significance* is determined, at least in part, by interests: whether or not a given height is *significantly* more than some other height depends in part on what interests things have (or would have). Here, the notion of *interest* is to be construed broadly, so as to include purposes, plans, desires, preferences, concerns, and intentions.

I might equally well have said that the predicate 'is tall for a cherry tree' is to be analysed as *has a lot more height than is the norm for a cherry tree*. 'A lot' is interest relative, just as 'significant' is.<sup>2</sup> If I say 'that is a lot of wine', then obviously whether what I say is true can depend on what my plans are—on whether I plan to serve the wine at a party of fifty people or at an intimate dinner for two. But 'significant' wears its interest relativity more prominently on its sleeve. In particular, and this is the source of Stanley's worry, 'significant' permits explicit

---

<sup>1</sup>See 'Shifting Sands: An Interest-Relative Theory of Vagueness' (Fara 2000) for my initial, and as yet most detailed development of the view.

<sup>2</sup>A predicate is interest relative when it expresses an interest-relative property.

relativization to an *agent* with purposes and interests; when something is said to be significant, we can appropriately ask to *whom*, or for *what*, it is significant.<sup>3</sup>

As I said in “Shifting Sands” (p.65)—and this actually turns out to be a pivotal point in the present debate—I do not want to be specific about what I mean by *analysis*. The worry was that in trying to be specific, I’d raise too many irrelevant side issues. Rather than be specific, I require only that *analysandum* and *analysans* have the same modal- and temporal-profile—that they be true in just the same worlds and at just the same times. I do not, for example, require that they have the either the same knowability- or believability-conditions (the same “epistemic profile”) or the same “constituents”, in case the Russellian or the Fregean view of propositions is correct.

To make things simple and concrete for the sake of argument, let’s just assume that the view is this: when I use the phrase ‘tall for a cherry tree’, it is to be analysed as meaning *significantly-to-me taller than is the norm for a cherry tree*. My proposal, in its barest outline, began with the following thought. Among whatever other interests I may have, I also have a standing interest in efficiency that causes me to avoid making discriminations that are too costly. By *making a discrimination* I mean making a *practical* discrimination—treating two things as unequal for the purposes at hand. It will be too costly to make a practical discrimination when the costs of doing so outweigh the benefits of doing so. Now let us suppose that my primary purpose is to choose a cherry tree for the yard. A discrimination between two cherry trees that are very similar in height will be very costly given my interest in efficiency. But the discrimination will be costlier still when I am actively considering the two trees as live options for my purpose. If the trees are similar enough in height, I’ll regard them as “the same for present purposes.” When I regard the trees as the same for present purposes, one tree will be significantly (to me) taller than the norm *just in case the other is*. For if that were not so, then I would not be treating them differentially for present purposes. Now, the thought continues, when I then go on actively to compare some *other* pair, perhaps equally similar in height, the discrimination cost of the new pair goes up, while the discrimination cost of the original pair goes down. Consequently, the change in what best satisfies my interests—as a result of the changing cost of discriminations between pairs, even in the absence of any change in the heights of the cherry trees—may result in a change in the *extension* of the interest-relative property being attributed to them.<sup>4</sup> In a slogan: the boundary between what is tall for a cherry tree and what is not is never where I am looking. The reason? In a slogan: the more “salient” a small difference, the more costly it is to recognise it from a practical point of view.

Let me dispense with slogans and be just a bit more concrete about some pertinent details of the view, particularly as it concerns gradable adjectives.<sup>5</sup> I take it to be a shared commitment that

<sup>3</sup>I leave it open that an “agent” may have interests even if it cannot in any straightforward sense act, think, or have intentions—such as a species (*e.g., the gorilla*), an area of inquiry (*e.g., cancer research*), or an abstract type (*e.g., punk rock*).

<sup>4</sup>I mistakenly gave the impression in “Shifting Sands” that it is the change in our interests that causes this change in extension, but that is not really a correct characterisation of the view I argued for there. It is rather a change in what best *satisfies* our interests—given our standing interest in efficiency—resulting from a change in what is being actively considered, that causes this extension change.

<sup>5</sup>By a *gradable adjective* I mean one that can meaningfully form comparatives—*e.g., ‘more fun’ and ‘taller’*. The categorisation is occurrence dependent. (‘More dead’ is in some, but not all, occurrences sensible.)

a gradable-adjective predication of the form ‘ $x$  is tall for a  $C$ ’ has a semantic analysis which can be represented schematically as follows:

(SA)  $x$  is tall for a  $C \equiv$  the height of  $x$  bears the relation  $R$  to  $f(C)$ .<sup>6</sup>

The idea is that in order to be tall, your height has to bear *some* relation to some height determined by the comparison class  $C$ . Presumably the relation  $R$  is at least as strong as the *greater-than* relation, but as far as this schematic representation goes, it is left open what function  $f$  is and what relation  $R$  is. It could be that  $R$  is not interest relative and that  $f$  is just a purely extensional averaging function. In fact, it is even left open that  $f(C)$  is not a height at all.

My view is that  $R$  is the interest-relative relation *is significantly greater than to/for a*, where  $a$  is some contextually given agent. I call  $f$  the “norm” function on comparison classes. This function will of course vary with the adjective in the sentence; when the adjective is ‘tall’  $f(C)$  is the normal *height* for  $C$ . But the adjective does not by itself completely determine what function  $f$  is—for there is more than one kind of norm. There are *e.g.* norms of expectation, idealistic norms, and various statistical norms. So  $f$  itself is also something that may vary with the context. In almost no context will  $f$  be a purely extensional averaging function.

So I do think that these vague gradable-adjective predicates are context dependent due to the contextual variance of  $R$  and  $f$  and also of  $C$  when the comparison class is not made explicit. In each context these predicates express an interest-relative property. It is the changing *extension* of these contextually determined interest-relative properties that leads to the boundary shifting which I appeal to in my solution to the sorites paradox; it is not the contextual variation in *which* property is expressed by these predicates that leads to this boundary shifting, as it is on the “contextualist” views of Kamp (1981), Raffman (1994, 1996) and Soames (1999).

## 2. Objection

Stanley pursues two lines of argument. First, there is the argument that the *epistemic* profile—in any context—is not what we would expect it to be if my interest-relative account were correct. In conflict (it is claimed) with my account, one can believe (or know) that the predicate applies to a thing without having any beliefs about me or my interests. You can believe my claim that a given cherry tree is tall without believing (or knowing) anything about *me*.

Second, there is the argument that the modal profile of ‘tall for a cherry tree’—in *any* context—is not what we would expect it to be if my interest-relative account were correct. In conflict with my account, the extension of this predicate in a possible state of affairs never depends (it is claimed) on what my interests are in that state of affairs. Its truth in a state of affairs does not even require my existence in that state of affairs. When I say that a given cherry tree is tall (for a cherry tree), the truth of what I say has nothing to do with me or my interests.

---

<sup>6</sup>To my knowledge, Bartsch & Vennemann (1972) were the first to propose a semantic analysis with this sort of structure. They took  $R$  to be the *greater-than* relation, and they took  $f$  to be the height-averaging function on the members of  $C$ . The resulting truth conditions have appeared in a number of works, and are to my mind entirely implausible due to the proposed extensionality of  $f(C)$  and also the weakness of the relation  $R$ .

Underlying these two arguments is a crucial premiss, which we could call the *aboutness claim*; the *reason* that the modal and epistemic profiles are not themselves interest relative is that, in any context, the claim ‘That cherry tree is tall’ does not express a proposition that is *about* me or my interests, in conflict (it is claimed) with my account.

### 3. Rebuttal Part A. Epistemic Involvement

Propositions with an interest-relative modal profile need neither be *about* nor *epistemically involve* interested agents, *contra* Stanley. This is true if propositions are coarse grained, *à la* David Lewis and Robert Stalnaker, with distinct ones always differing in truth conditions. And it is also true if propositions are finer grained, *à la* Frege or Russell, with some distinct ones having the same truth conditions but different “constituents”. To see that this holds on both of these competing views of the identity conditions of propositions let us suppose that propositions  $P$  and  $Q$  have the same modal profile; they are necessarily equivalent propositions.

First, if the Lewis–Stalnaker coarse-grained view of propositions is right, then  $P = Q$ . A coarse-grainer will think that there is more than one way for me to know this proposition. Knowing a certain something about  $X$  might suffice for me to know the proposition, and also knowing a certain (non-equivalent) something about  $Y$  might suffice for me to know the proposition. This  $X$ -knowledge and this  $Y$ -knowledge are each sufficient, although neither is necessary. An example: I could know the necessary proposition by knowing that this proof concerning  $\pi$ , written on this board, is correct. Another way could be for me to know that *that* proof concerning the number two, written on *that* board, is correct.

Second, if one of the Russellian or Fregean finer-grained views is right, then  $P$  and  $Q$  may be different predications with different constituents and accordingly may be *about* different things and have different knowability conditions. If a Russellian, for example, proposes an “analysis” of a certain property by offering up a distinct property claimed to have the same modal profile, she is not thereby committed to propositions involving the one property being about, or epistemically involving, just the same things that propositions involving the other property are.

I have offered up an analysis of gradable-adjective predications in just this spirit. Since my claim about analysis is only a claim about modal profile I have absolutely no commitments regarding *aboutness* or *epistemic involvement*, no matter which view of propositions is the correct one.

I suspect that some Russellian conception of propositions is correct. The view would be that propositions have as constituents individuals and properties and functions of many types, namely, those serving as the denotations of pronounced and unpronounced syntactically simple elements.

Let me therefore explain, assuming this particular Russellian picture, why my view does not have interested agents as constituents of gradable-adjective propositions. The view hinges on a key hypothesis: just as comparative constructions with gradable-adjective phrases have their own morpheme (‘-er’), and superlatives and equatives their own (‘-est’ and ‘as ... as’), seemingly-

bare *positive* constructions have their own—unpronounced—morpheme.<sup>7</sup> This allows for an appealingly unified account of the semantics of gradable adjectives as they occur in their various constructions, and of the syntax of these various constructions. In “Shifting Sands” I borrowed and adapted Christopher Kennedy’s implementation of this hypothesis, according to which gradable adjectives univocally denote functions from objects to degrees on an associated scale.<sup>8</sup> My adaptation of Kennedy’s proposal includes the following theses.

Whereas the semantic analysis of ‘ $x$  is tall for a  $C$ ’ could be schematised as follows:

(SA)  $x$  is tall for a  $C \equiv$  the height of  $x$  bears relation  $R$  to  $f(C)$ ,

its syntactic structure should be represented like this:

(SS) [ $x$  [[ $\emptyset$  tall]  $C$ ]].

Here  $\emptyset$  is the unpronounced positive morpheme (earlier, Kennedy (1997) called it the “absolute morpheme”); *tall* denotes a “measure function”—specifically, a function that measures things’ heights (or, perhaps, degrees of tallness); and  $C$  is a comparison class. So what type of function or property must the positive morpheme  $\emptyset$  be in order to achieve the semantic interpretation (SA)? Given the order of composition dictated by the syntactic structure (SS) it will be a high-type function having measure functions for its domain and functions from comparison classes to properties of individuals for its range. On my interest-relative theory, it is a function  $\emptyset$  such that  $(\emptyset(G))(C)$  is a property that is true of a thing  $x$  just in case  $G(x)$ ,  $x$ ’s amount of  $G$ -ness, is significantly (to  $a$ ) greater than the typical (“norm”) amount of  $G$ -ness for a  $C$ . Which function precisely this is will depend on which of the various norms, and what agent  $a$ , is operative in the context.

When Stanley says (2003: 278) that according to this view the positive morpheme denotes the *significantly-greater-than* relation which requires as an implicit argument an agent with interests, and that the view therefore requires there to be interested agents as constituents of propositions expressed using gradable-adjective predications, he glosses over the pertinent aspect of my view. The positive morpheme does not denote a relational expression, but rather the high-type function just described. So in gradable-adjective predications, there is no relational expression—pronounced or not—that takes an interested agent as argument, hence no interested agent occurring as a constituent of these propositions. My view does not, therefore, have the perhaps objectionable consequence that vague propositions are always about, or epistemically involve, interested agents.

#### 4. Rebuttal Part B. Modal Involvement

In order to convince us that vague propositions do not have an interest-relative modal profile, Stanley has us consider John’s utterance of the following sentence.

(1) Mount Everest is tall for a mountain.

<sup>7</sup>The hypothesis can be found in a number of works on the semantics of adjectives, e.g., in (Cresswell 1976), (von Stechow 1984), and (Kennedy 1997).

<sup>8</sup>A measure-function analysis can also be found in Bartsch & Vennemann (1972)

He writes:

According to [Fara], this occurrence of (1) expresses the proposition that Mount Everest is significantly for John taller than the typical mountain. *But what (1) expresses could still be true, even if John had never existed* (and hence didn't have interests or purposes). (2003, p. 278, emphasis added)

Notice how Stanley's claim about the modal profile of John's utterance involves an embedding within a 'might'-counterfactual. Notice also how he is careful not to make his claim about modal profile by asserting the following different 'might' counterfactual:

(2) Mount Everest could still have been tall even if John hadn't existed.

This is a modal claim about Mount Everest, not a claim about the modal profile of any proposition. For this reason, Stanley instead uses the locution "what (1) expresses" in order to designate the proposition expressed by (1) and to say something about *it*. But, I claim, embedding the "what-(1)-expresses" locution within the 'might'-counterfactual as he does exploits our correct judgement that (2) is true in order to get us to conclude—wrongly, I claim—that the modal profile of (1) is not interest relative.

### **Modal Sentences and Modal Profiles**

One can distinguish the question of the truth value of a modal claim, such as 'that tree might have been tall', from the question of the modal profile of its embedded clause, 'that tree is tall'. I here assume a standard possible-worlds framework which identifies the modal profile of a sentence with the set of possible worlds in which it is true.<sup>9</sup> Within this framework possibility or necessity is truth in some or all (respectively) possible worlds—usually with the requirement that those worlds meet some particular condition. Once in this framework, we may say that the modal claim is true when the modal profile of its embedded clause contains at least one world meeting the required condition.

One reason for making this distinction is that the truth value of the modal claim might be affected by ambiguities where there are no ambiguities in the clause embedded in it. This difference can create distractions when it comes to assessing the truth of the modal claim, that aren't there when assessing the modal profile of the claim's embedded clause. These could be scope ambiguities: 'The inventor of the zip must have invented the zip' contains a scope ambiguity that can affect its truth value while 'the inventor of the zip invented the zip' contains no scope ambiguity. And, what is more relevant for us here, these could be ambiguities resulting from the effect of presuppositions on the interpretation of modal auxiliaries. For example, if we are presupposing that women are barred from professional baseball, as they in fact are, then 'I could have been a shortstop for the Mets' is false, but still the modal profile of 'I am a shortstop for the Mets' is

---

<sup>9</sup>Thus "modal profiles" are what are considered to be propositions by those, like Robert Stalnaker and David Lewis, who favour a coarse-grained conception of propositions, as contrasted with Russell's more fine-grained conception or Frege's even finer-grained conception.

non-empty. It merely lacks a world of the required kind, namely, one that's compatible with our presuppositions.

Even when we pose questions clearly intended to be ones about modal profile, the use of modal locutions may taint our judgements. For example, suppose I make the absurd claim that I *am* a shortstop for the Mets. If we then go on to ask whether what I said *might* have been true, we may be inclined to say 'no' since the interpretation of 'might' in our question about profile can be affected by the very same presuppositions as before about women and baseball. It is precisely this sort of distortion of our judgements concerning modal profile that makes Stanley's argument concerning the modal profile of interest-relative propositions so enticing.

Given that claims about modal profile, when stated counterfactually, are liable to be interpreted like their object-language counterparts, the fact that "what (1) says" could have been true even if there had been no interests, does not render it obvious that (1) is not interest relative. In the service of arguing that this distortion is artificially bolstering Stanley's case against my view, I will show that there are uncontroversially interest-relative sentences that embed in counterfactual conditionals in the way that I'm appealing to.

### Counterfactuals and Modal Profiles

Suppose some anthropologists are talking about a cave with an interior wall that seems to them to be unnaturally smooth. They are debating whether its smoothness is evidence of humans' having been there. One of them who thinks that a cave wall could be that smooth only if some humans had smoothed it might well say (and truly so, at that):

- (3) If no people had ever existed it would be very surprising for this cave wall to be so smooth.<sup>10</sup>

'Surprising' is an interest-relative word. This is why the modal profile of an utterance of (4) below contains no world in which there are no interests (expectations, *etc*). In short, no interests, no surprise.

- (4) The smoothness of this cave wall is very surprising.

So how can this be the case? How can it be that the conditional in (3) is true, even though there are possible worlds in which its antecedent is true, but none of them are among the worlds in the modal profile of the consequent?

The reason is that the evaluation of counterfactuals often involves holding certain facts fixed, even when those facts *would not have obtained* had the antecedent of the conditional been true. In the case of (3) what is being held fixed is which sorts of things are surprising. The apparent effect is the "rigidification" on what peoples interests are, but without the involvement of any "actuality" operator. Another example is this. You say to me, "If I had held an interest-relative theory of

---

<sup>10</sup>There is a related but truth-conditionally different counterfactual in the vicinity in which 'surprise' takes wider scope: "It would be surprising if this cave wall could be so smooth if no people had ever existed". The difference is that between  $\sim P \Box \rightarrow S : C$  and  $S : (\sim P \Diamond \rightarrow C)$ .

vagueness, then we'd have similar theories.” In this case, we are holding fixed the fact that I hold an interest-relative theory of vagueness, even though I distrust your views as a philosopher so much that I would have held a different view had yours been interest relative, in order to be different from you.<sup>11</sup>

## 5. Conclusion

Stanley claims that vague propositions would have to be *about* and *epistemically involve* agents with interests if the interest-relative theory were correct. We saw that this allegedly bad consequence would not follow no matter which view of propositions were the correct one.

Stanley claims that vague propositions would have to have an interest-relative modal profile if my view were correct. I agree, but deny that this is problematic. Stanley's argument that vague propositions are not interest relative is dubious. Counterfactual claims about sentences and the propositions they express can and do get conflated with their object-language counterparts. Moreover, a non-modal version of this objection, reformulated so as to avoid my response:

- (5) The proposition expressed by (1) is not true in any world in which there are no interests,

does not constitute an argument against my theory, but merely a flat-out denial of it.<sup>12</sup>

## References

- Bartsch, Renate & Vennemann, Theo (1972), 'The grammar of relative adjectives and comparison', *Linguistische Berichte* **20**: 19–32.
- Cresswell, M.J. (1976), 'The Semantics of Degree', in B. Partee, ed., *Montague Grammar*, Academic Press, New York, pp. 261–292.
- Fara, Delia Graff (2000), 'Shifting Sands: An Interest-Relative Theory of Vagueness', *Philosophical Topics* **28**(1): 45–81. Published under the name 'Delia Graff'.
- Kamp, Hans (1981), 'The Paradox of the Heap', in U. Mönnich, ed., *Aspects of Philosophical Logic*, D. Reidel, Dordrecht, pp. 225–277.
- Kennedy, Christopher (1997), *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*, PhD thesis, University of California, Santa Cruz. Published (1999) by Garland Press, New York.
- Raffman, Diana (1994), 'Vagueness without Paradox', *Philosophical Review* **103**(1): 41–74.
- Raffman, Diana (1996), 'Vagueness and Context-Relativity', *Philosophical Studies* **81**: 175–92.
- Soames, Scott (1999), *Understanding Truth*, Oxford University Press, New York.
- Stanley, Jason (2003), 'Context, Interest Relativity and the Sorites', *Analysis* **63**(4): 269–280.
- Stanley, Jason (2005), *Knowledge and Practical Interests*, Oxford University Press, Oxford.
- von Stechow, Arnim (1984), 'Comparing Semantic Theories of Comparison', *Journal of Semantics* **3**: 1–77.

---

<sup>11</sup>It's because of cases like these that *strengthening* fails for counterfactuals.

<sup>12</sup>Many thanks for comments to Gilbert Harman and Katrina Przyjemski and for discussion to Michael Fara, Christopher Kennedy, Jeffrey King, Jason Stanley, Timothy Williamson, and the audience at the 11<sup>th</sup> Annual Oxford Philosophy Graduate Conference, Oxford University, November 2007.