This is a preprint of an article whose final and definitive form will be published in *The Australasian Journal of Philosophy*, 2008. *The Australasian Journal of Philosophy* is available online at: http://journalsonline.tandf.co.uk/.

Gradable Adjectives: A Defense of Pluralism

Keith DeRose
Yale University

1. The "For an F" Myth (or Implicit Reference Class Theory) Regarding Gradable Adjectives and My Motive for Resisting It

When you've applied a gradable adjective to something – say you've claimed that "Sally is tall" – and you are asked what you meant by the claim, one quite common type of answer is to cite a comparison class (or reference class) relative to which your use of the gradable adjective is to be understood: "What I meant was that Sally is tall for a 10-year-old." And if you've first claimed that, say, "Sally is tall," but in a later conversation say that "Sally is not tall," and your consistency is challenged – "But you said earlier that Sally is tall" – it's quite common again to clarify the meaning of your statements by appeal to comparison classes: "All I was claiming then was that she is tall for a 10-year-old; what I'm saying now is that she isn't tall for a basketball player." Focusing too myopically on cases like that might lead one to accept what I call the "'for an F' myth," or as it has been more neutrally labeled (e.g., [Richard 1995]), the "Implicit Reference Class Theory" (IRT), of gradable adjectives. According to this myth/theory, where 'A' is a gradable adjective, claims of the form "X is A" quite generally mean that "X is A for an F", where F is some comparison class that is supplied by context. So, what context does in fixing the content of a particular use of a gradable adjective is supply the comparison class relative to which the claim is made.

The bulk of this paper will be devoted to attacking this "for an F" myth, and replacing it with what we can call a "pluralist" approach to gradable adjectives (adjectives like "tall", "short", "big", "rich"), according to which gradable adjectives can be governed by a variety of different types of standards, one, but only one, of which is the group-indexed standards utilized by the Implicit Reference Class Theory. The attack begins in section 2, below. But first, I will

briefly explain the reason for my interest in this issue – beyond the intrinsic value of getting the semantics of gradable adjectives right.

I don't know how widely the "for an F" myth is actually accepted – particularly by those who have spent much time actually thinking about the semantics of gradable adjectives. My target here is not the literature devoted to adjectives in linguistics and philosophy journals, and I will confine my rather brief remarks about the extent to which that literature seems complicit in the myth to a footnote.[1] But I do know from personal experience that the myth is quite widely assumed by many of those who (and especially by philosophers who) don't work specifically on gradable adjectives, but who do use what they think they know about how gradable adjectives work in thinking and arguing about other cases of alleged context-sensitivity. And I suspect that the assumption of the myth is at least partly responsible for creating an environment in which there is a good deal of resistance among philosophers to various "contextualist" positions in philosophy.

---

[1] First, it's important to note that I am not using the label "IRT" to cover just any treatment of gradable adjectives that somehow or other makes use of the idea of implicit comparison classes. Much of the most important literature on these adjectives uses syntactic evidence to yield conclusions about the role that implicit comparison classes play in at least some uses of gradable adjectives. Such conclusions do not get you to what I'm calling "IRT," according to which, where 'A' is a gradable adjective, "X is A" quite generally means that X is A for an F.

So, how complicit in the myth is the literature that is devoted to gradable adjectives and that uses the notion of comparison classes? Here I'll confine myself to one very prominent example of the literature in question: [Ludlow 1989] – though much of what I say about it applies to other work as well. First, there seems to be no hint at all here of what in section 3 I will call the "strong" form of IRT, according to which it's not only the case that gradable adjectives are quite generally governed by what I'll call group-indexed standards, but also that supplying the relevant comparison class is all context ever does in fixing the content of a use of a gradable adjective.

However, there does seem to be suggestion of what I call the "weak" form of IRT, which claims only that "X is A" quite generally means that "X is A for an F," without also insisting that all context does in fixing the content of such a claim is supply a comparison class. Ludlow's paper opens with these words: "A common view of attributive adjectives like 'tall', 'fast', 'large', and 'heavy' is that they express relations between objects and either some comparison class (hereafter c-class) or an attribute" [1989: 519]. This "common view," then, is not presented as one on which such adjectives sometimes or often express relations between objects and c-classes, but one on which such relations simply are what the adjectives express. Now, it also is not explicitly specified that on the "common view" the adjectives always express such relations to c-classes, but in this context, the simple statement that such adjectives express such relations (as opposed to, say, that they often express such relations) strongly suggests something at least approximating that they always express such relations. (Perhaps the best reading is something along the lines that the adjectives at least normally express such relations.) So, that "common view" seems to be at least very close to what I call the weak form of IRT. After a little more explanation, Ludlow continues: "If we adopt this view, an interesting question arises..." [1989: 519], and he proceeds to address that question, and others, and the "common view" is never withdrawn. It seems that the "common view" is being assumed – and perhaps even to some extent defended, since Ludlow seeks to show how the questions that arise can be answered.

However, Ludlow's focus is not on defending the "common view," but on other conclusions. He seems willing to assume the "common view" and to work within it to get at the questions that are his focus. But the supposition that many uses of gradable adjectives are not governed by group-indexed standards would seem to do little to upset the conclusions that are Ludlow's real focus. Thus, those who assume IRT, even in its weak form, would seem to be making a mistake if they appealed to what's shown in [Ludlow 1989] in support of their assumption; Ludlow is setting his sights on quite different targets.

In several areas of philosophy, some recent attempts to solve various problems have crucially involved claims to the effect that certain key terms involved in the statement of the problems are context-sensitive. One prominent example of this phenomenon is the "hidden indexical" account of belief ascriptions. Though I have some interest in defending several such contextualist positions, including hidden indexical accounts of belief ascriptions, my own investment in such approaches is heaviest in the case of another prominent example: "contextualist" accounts of knowledge attributions. So that's the example I'll explicitly address in what follows. These various appeals to context-sensitivity have generated a good deal of controversy in recent years, and have active critics and defenders. But my focus right now is not on the explicit arguments given against these appeals to context-sensitivity, but rather on what I think is a widely assumed general picture of how context-sensitivity works in natural language that I suspect is working in the background. This general picture creates a hostile environment for the acceptance of contextualist theories, and is encouraged by the assumption of IRT in the area of gradable adjectives.

On the picture in question, context-sensitivity is not only quite limited in natural language, but also, where it does exist, it works very neatly. First, the work context does in fixing the content of particular uses of context-sensitive language is generally done by supplying, in the form of handy little phrases that the speakers themselves will supply if you ask them what they mean, specifications of the particular value a context-sensitive term takes on the use in question. Second, there is a very limited menu of the types of such items that context can supply, and hence of the varieties of context-sensitivity in natural language. There are the true indexicals, where the time (for "now"), place (for "here") or agent (for "I") of utterance is straightforwardly supplied by context (and can be specified by the speaker, if asked, by means of handy little phrases); then there are gradable adjectives, where context simply supplies the relevant comparison class, again specifiable by means of handy little phrases; and then perhaps a few other varieties of context-sensitivity – like pronouns, where a particular referent is supplied by context.

Against this very neat background, the workings of "knows," for example, can appear problematically messy. One can begin to wonder whether our language would really contain context-sensitive expressions without also providing neat and precise ways of indicating what particular content they bear on particular occasions.

Now, speakers do use various devices by which they at least arguably clarify what they mean by saying that someone "knows."[2] But none of these "clarifications" of "knows" work by supplying a comparison class. And this is already trouble in the minds of some who are in the grip of the "for an F" myth about gradable adjectives. For many are led to accept the second feature of the general picture, thinking that there is a very limited menu of items that can be supplied by context in fixing the content of a context-sensitive expressions: either the time, place, or agent of utterance, or some comparison class, or perhaps a very few other types of things, which don't include anything like epistemic standards. And these folk will not take kindly to the suggestion that there's a context-sensitive term for which what context supplies in fixing content is something not on their highly limited menu.

Beyond the fact that these "clarifications" for "knows" in natural English don't make various approved short-lists of items that context can supply in fixing the content of a context-sensitive expression, they also seem to be quite a bit more messy and imprecise than are the things on the approved list. What contextualists appeal to in explaining how "knows" works – varying standards for how strong an epistemic position one must be in to count as "knowing," or some similar mess to roughly the same effect – seem more disheveled than different comparison

---

[2] A few examples:
"All I was claiming was that I know it...
-quite well
-beyond any reasonable doubt
-by ordinary standards
-by any reasonable standard
-with a high degree of precision
...I never meant to be claiming to know it...
-for certain
-with absolute certainty
-beyond all possible doubt
-perfectly
-as God would know it."
These seem to be perfectly fine pieces of English, not cases of mere "philosopher-speak" – though philosophers, too, have occasion to say such things. The challenger might not like these responses, and might have objections to them, but the responder doesn't seem to be abusing the English language in using these as devices of clarification. For many more such examples, see [Ludlow 2005]. It's important that these are not just modifiers that "S knows that P" can take, but function as devices of clarification: In a context in which one's meaning or consistency are questioned, these devices at least arguably can function to clarify the content of one's claim. By contrast, while "quite well," "with a high degree of precision," "as God would play it," and "perfectly" can be used to fill in the blank in "She played the piano piece _____," the resulting sentences don't seem to function as devices of clarification.

classes. If someone explains that by "knows" they meant "knows even by quite demanding standards" (or if they give any of the "clarifications" listed in note 2, above, or in the paper by Ludlow cited in that note), the listener doesn't feel they've been given a very precise specification of the speaker's meaning. By contrast, however misleading this sense is (for reasons we'll see in section 2, below), it's easy to think that "tall for a 10-year-old girl" gives a quite precise meaning for "tall"; now you've been told exactly what was meant, and what was meant was quite precise. Thus, even someone who doesn't presume that the approved list is all there is might be <u>especially</u> worried about supposing that "knows" is context-sensitive in the way contextualists claim. So, again, while in the grip of the myth, it's easy to wonder how plausible it is to suppose that our language would contain a context-sensitive expression without also containing various devices for neatly and precisely fixing its content on particular occasions. For, while under the spell of the myth, it can seem that we have just such devices in the case of gradable adjectives.

So it's important to see that things really aren't so neat and precise, even when it comes to those parts of the language that are uncontroversially context-sensitive. In fact, though this isn't our present topic, even the true indexicals (or "core" indexicals) can be quite messy.[3] Our

[3] This is most vividly displayed by "here" – though similar problems can arise with the temporal extent of "now", and, as very strange imagined cases can show, with the issue of just how far out, temporally and/or spatially, "I" reaches. "Here" is used to refer to the location of utterance, all right. (Well, this already seems questionable, because of cases like those in which a speaker says, while pointing to a spot on a road map, "The exit we want is here." How and whether to handle such cases within a general theory of the workings of "here" can get to be fairly complicated business.) But what are the boundaries of the region designated by a given use of "here"? How far out does a given use of "here" reach? Sometimes we have simple little phrases that can explain fairly precisely just what we mean: To echo an example that Stephen Schiffer uses, we might mean "in the building", "on campus", "in the States", etc. – though on some occasions it might be unclear exactly which we mean [Schiffer 1995: 112-113]. Often, when I use such a simple phrase, I feel that I could just about as well have used another – a bit broader or narrower – instead: Instead of "in the building," I could just as well have said I meant "somewhere in the philosophy department," when I said that Larry wasn't "here". It at least feels as if, while I am somehow giving my permission to let my earlier use of "here" to be understood in the specified quite precise way, what I meant at the time I said it was vague enough, even in my own mind, that a different precisification would have been just as faithful to what I meant at the time of utterance. Schiffer presents a case in which it isn't clear which of many possible specifications of the meaning of "here" is right, and advocates a supervaluational approach to such situations. Schiffer does seem to assume, though, that each of the candidate specifications for the use of "here" in question can be given by means of a handy little phrase. But there are cases in which such handy phrases don't just fail to precisely provide <u>the</u> meaning in context of "here," but fail to even specify the candidate meanings among which the given use is indeterminate. Sometimes handy little phrases simply fail us, as we can see by moving to a different example. "Is Michael still here?", a speaker might ask, while looking about for Michael, in a conversation taking place

present topic is the messiness of gradable adjectives – which are at least fairly uncontroversially context-sensitive. So we now begin the attack on the "for an F" myth, partly in the hopes of upsetting the neat general picture of context-sensitivity seemingly assumed by many.

Like many myths, the "for an F" myth has some basis in fact. Uses of "tall," for example, can be governed by a variety of different standards. And here is what the myth gets right: Some of these standards are indeed of the "for an F" variety; we can call these "group-indexed" standards. But, as I will argue, beginning in section 4, such group-indexed standards are only one out of several types of standards that can govern uses of "tall" and the like: Uses of gradable adjectives are often instead governed by one or the other of the very different types of standards that I will call "tape measure" and "group-relative" standards. So, "X is tall" does not generally mean that "X is tall for an F." But before getting to that, we will see in section 2 that even when a use of a gradable adjective is governed by some kind of group-indexed standard, there can be a difference in content without there being a difference in what the comparison class is, so that fixing the reference class does not suffice to fix the content, and section 3 will discuss the importance of that to the concerns that motivate this paper.

---

at one of the "smokers" held in giant ballrooms at APA conventions. Well, obviously, Michael isn't right here, within just a few feet of us. What could the speaker mean? "Here in the ballroom," perhaps? That's the only simple and fairly precise phrase that seems available. But, no, Michael himself was taking part in our conversation just a minute or so ago, before he drifted a few feet over to say a few words to someone else who was standing nearby. Now he's apparently drifted even further away, but we're all quite certain he's still in the ballroom, and the speaker wouldn't be asking that. But what, then? Given the way the speaker casts her eyes about as she asks the question, she seems to mean, however vaguely, to be designating a region of the ballroom in which there's a relatively high percentage of people known to us. (Has Michael just moved a little further away, to a spot around here, or has he perhaps slipped out of "the region" to go to the bar and get a new beer or to talk to a whole new group of people in some other area of the ballroom?) But if you were to ask the speaker to specify exactly what area she meant to be designating by her use of "here," she'll just laugh and explain that she didn't mean anything very exact: "You know, right around here," she might say, perhaps making vague, sweeping gestures with her hands. No simple phrase is available, and, for certain spots at some distance, but not a great distance, from the speaker, even the speaker has no idea whether they would or would not count as being in the region she meant to be designating with her use of "here". What's more, there seems to be no way, at least by means of simple little phrases, to precisely identify even various candidate meanings that the speaker's use of "here" is indeterminate among.

2. Varieties of Group-Indexed Standards

Here we are following the lead of Delia Graff Fara, who has argued that "not every variation in standards of use for vague expressions can be attributable to a variation in implicit comparison classes" [2000: 56].

"X is tall," I admit, can mean something like that "X is tall for a 10-year-old." But how tall for a 10-year-old must X be to count as being tall for a 10-year-old? Taller than most 10-year-olds? Taller than the vast majority of 10-year-olds? Upper 20% of 10-year-olds? Upper 10%? Upper 2%? Many of our uses of "tall" seem too vague to be susceptible to the numerical precision displayed by some of the above options; we often have nothing so precise in mind. But our at least apparently more vague uses of "tall" vaguely approximate different of these options on different occasions,[4] generating what we can call "serious context-sensitivity": Sometimes

[4] Standards that can be specified simply by citing the threshold for what percentage of the Fs objects must be A-er than to count as being "A" are a sub-type of what we can call "roster standards." Roster standards are functions from "rosters" – specifications of how A each of the Fs is – to what, starting in section 4 below, we will call "tape-measure standards." (The idea is that if a standard is a roster standard, then the only information one needs to generate a tape-measure standard from it is the information from the relevant "roster" – specifications of how A each of the Fs is.) So, for instance, the group-indexed, taller than at least 98% of 10-year-old boys, is a roster standard, because, to generate a tape-measure standard for tallness from it, the only information one needs is how tall the various 10-year-old boys are. (Since I have little idea how tall 10 year old boys tend to be, I have little idea what tape measure standard is generated by the roster standard in question.) Roster standards need not be matters of simple thresholds of percentages, but those that work by simple percentages in that way are the only ones that seem easy to specify and explain. In fact, it's possible that we resort to simple percentages to explain our standards simply because those are the only easily specifiable standards in the vicinity, and that many of our "roster" uses of adjectives are in fact more complex, despite the simple way we would specify them if asked. I say, "Ralph is rich," and I'm asked what I mean. "Rich for an American," I reply. "How rich for an American?," I'm asked. Meaning to employ fairly demanding standards, I reply, "Oh, in the upper 5% of Americans." But how would my use of "rich" apply to a counterfactual situation in which the distribution of wealth in America were even more unequal than it is? What if well over 99% of the America's wealth were held by the upper 2% of Americans, and the remaining 98% of Americans were all desperately poor, and all very close to one another in how wealthy they were? Is my actual use of "rich" such that it would count someone as rich for an American if they were in the upper 5%, but not the upper 2%, of Americans under those circumstances? If Ralph were in that position – in the upper 5%, but only marginally wealthier than the very poorest Americans, and, like almost everyone else, much poorer than the upper 2% -- it seems that one wouldn't have to be employing particularly demanding standards to count him as not being rich. But he would be in the upper 5%, and so would count as "rich" if the standard really were simply that of being in the upper 5% -- which is the standard I had specified, thinking that I was employing a very demanding standard. Perhaps, despite what I'd say if asked what I mean, my use of "rich" would be more accurately captured by a more subtle roster standard (a more subtle function from rosters to tape measure standards), but, since my listeners and I would naturally assume a more normal and even distribution of wealth than exists in our imagined situation, there's no

merely being in the upper quarter of 10-year-olds in height clearly or definitely counts as being "tall" for a 10-year-old, and sometimes it quite clearly or definitely doesn't so count. So there is context-variability that does somehow get resolved on occasions, though it isn't resolved simply by providing the relevant reference class, as has been noticed, for instance, by John Hawthorne: "Note that, even having a fixed reference class, there may be context-dependence generated by the need for a standard of application – some kind of threshold on a scale – relative to that reference class. (To be tall do you have to be taller than 80% of the reference class? 85%?)" [2004: 53-54].

That and more has also been noticed earlier by Fara, who presents several different types of cases of variations in the standards governing a gradable adjective without a variation in what seems to be the relevant comparison class. Some of Fara's cases suggest that we sometimes use group-indexed standards for gradable adjectives that don't operate in the way that, like Hawthorne, I focus on above (by some threshold of the percentage of items in the reference class that something must be A-er than to count as being A), but involve interestingly different types of group-indexed standards or norms.[5] For instance, consider Fara's Fido in this passage:

> Suppose that Fido is fourteen years old and Rover is twenty years old. Someone who says that Rover is old for a dog may be making a remark about his extreme longevity, while someone who says that Fido is old for a dog may be merely remarking that he is in his old age. This is to be explained, on my view, by a variation in norm. Rover, the twenty-year-old, has significantly more age than is the norm for a dog to attain; while Fido, the fourteen-year-old, has significantly more age than some different kind of norm for a dog, one that's much harder to articulate, but which perhaps concerns the peak age of good health. [2000: 66]

---

harm done in specifying what I mean by simply saying "upper 5%," not worrying about whom would be counted as "rich" under various extreme counterfactual conditions.

[5] By "interestingly different" here, I specifically mean that Fara's cases seem to involve group-indexed standards that are not what, in note 4, above, I call "roster standards": One cannot determine a tape-measure standard for oldness for dogs from the group-relative standards Fara is pointing out, using only the "roster" information of how old the various dogs are.

Fara is here writing above about uses of an expression that are explicitly relativized to a comparison class – "old for a dog". But, as Fara realizes, simple uses of "old" – "Fido is old", "Rover is old" – can inherit this variation in the types of standards or norms that can govern their meaning, even when the implicit comparison class – in this case, that of dogs – is held constant.

3. The Downfall of the Strong Version of IRT

In the above section, we saw cases which show that fixing the reference class does not fix the content of a given use of a gradable adjective.[6] This already refutes IRT, at least in the strong version of the theory that I designated – one according to which all context does in fixing the content of a gradable adjective is fix the reference class, so that all differences in the content of a given gradable adjective will be due to differences in references classes. But some weakened version of IRT – one that allows that there is a variety of different group-indexed standards that can be used to govern whether something counts as being A, even where the reference class is held fixed, but continues to insist that all uses of "X is A" are governed by group-indexed standards – can survive the examples we've seen so far.

Starting in the next section, we'll discover that even the weak version of IRT is a myth, when we see perfectly good uses of gradable adjectives that aren't governed by group-indexed standards at all. But before moving on to that, it's worth noting how the downfall of just the strong version of IRT does much to address the worries that are here motivating my investigation of gradable adjectives.

So, first, as I noted, some in the grips of the myth can become addicted to a highly limited menu of items that context can supply in fixing the content of context-sensitive language (time, place or agent of utterance, or some reference class), and will be suspicious of the contextualist's claim that varying standards for knowledge can be supplied by context, since those varying epistemic standards don't appear on their short-list of approved items. But if two uses of, say, "Sally is tall" differ in content from one context to another even while there is no difference in the relevant comparison class, as we've clearly seen can happen, then context must

[6] Here I am not assuming that content cannot be vague. Of course, it can be. But the uses of the adjectives that we looked at have (fairly precise) contents that could not have been fixed simply by specifying the comparison classes relative to which the uses are to be understood.

be supplying something here that is not on the various approved short-lists – whether it's in addition to or instead of a reference class.[7] So, for example, if we have two uses of "Sally is tall" where the comparison class is adult females, but in one case Sally only has to be taller than most adult females to count as tall, but in another she has to be in about the upper 2% of adult females in height to so count, then the contexts of these two uses must be supplying something beyond the comparison class of adult females in fixing the contents of the claims – they are also in this case each supplying some standard for what's to count as a tall adult female.

Second, there's a worry about precision: Even those who don't presume they have a complete short-list of items that context can supply in fixing content can, while under the spell of the myth, think that whenever our language has context-sensitive terms, then either the values that context supplies are straightforwardly and automatically supplied quite precisely by context – by the time, place, or agent of utterance, in the case of core indexicals – or else the language also supplies us with neat devices by which we can quite easily and quite precisely state exactly what we meant by a given use of a context-sensitive term – "for an F" constructions, in the case of gradable adjectives. In either case, whenever speakers sensibly use a context-sensitive term, they know which value governs their use, and if you ask them what they meant by their use, they will be able to tell you quite precisely by, say, stating their location (in the case of "here"), or stating the comparison class relevant to their use of a gradable adjective: "What do you mean Yao is tall?" "I mean he's tall for an NBA center." By contrast, in the case of knowledge attributions, when we ask a speaker what exactly they meant, then, if they're able to answer at all, their answers are likely to seem quite imprecise: "I meant I know it with a reasonable degree of certainty." And, as I noted, one can then wonder whether our language would really contain a context-sensitive term without also providing neat devices by which one could precisely state what one means by a given use of the term.

But we can now see that when speakers clarify that by "A", they meant "A for an F", their clarification is in reality not precise at all, for "A for an F" is itself very imprecise, and, what's more, can itself mean very different things on different occasions.

[7] This conclusion is also shown by our brief look at the workings of "here" in note 3, above. The region of the ballroom vaguely designated by our speaker's use of "here," without the benefit of any handy little phrases to specify it, would also seem to be something supplied by context, though it's not the kind of thing that seems to make the assumed short lists.

4. Tape-Measure Standards

There are also uses of gradable adjectives that seem not to be governed, even vaguely, by any type of group-indexed standard at all. I'll point out two other types of standards, investigating "tape-measure" standards in this section, and "group-relative" standards in section 6, below.

As I noted in section 1, sometimes when we seek to clarify or explain what we mean by a given use of a gradable adjective, it's quite natural for us to do so by explaining what comparison class our use is relative to. But on other occasions, it would be extremely unnatural for us to do that – and what we would naturally say points to a different type of standard.

Consider a movie director setting up the background for a key scene. "I need something tall over there on the left, to balance the shot – maybe a tree, or a streetlight or something. Get me something tall!" When asked what she means by "tall", the natural answer for her to give might well make no reference at all to any comparison class, but might rather simply cite what we may call a "tape-measure" standard: "What I mean is something about 14 to 16 feet tall." It might not only be natural not to cite any reference class, but it might be hard to say what comparison class could possibly be instead cited. The director has mentioned trees and streetlights, but she's not looking for things that are tall for a tree or tall for a streetlight; in fact, what she seeks might well be quite short for a tree or for a streetlight. Is there some implicit reference class that the director hasn't mentioned? Tall for a thing on this set? What she seeks might not be something tall for a thing on this set. Suppose the set already contains several 25-plus foot-tall houses in the center of the set. The only other item is an about 15-foot-tall tree on the right side. And now the director seeks to put something over on the left side that's about as tall as that tree. So what she seeks isn't tall at all for a thing on this set – in fact, it would be about tied for the least tall thing on the set. So what is the reference class relative to which the director is saying she needs something "tall"? Maybe there are no good candidates. ("Tall for a thing"?!) It seems that such a use is simply governed by a tape-measure standard, and is not made relative to any comparison class at all. And when a suitable item finally is located, the "tall" in "Yes, that's tall" seems to be governed by a tape-measure standard.

Since I've found that this is the point at which some die-hard IRT backers dig in, it's worth following up with some further explanation for why I think such uses as my director's aren't governed by implicit reference classes. Perhaps I was too quick above in dismissing such

a general reference class as "things." Indeed, in his classic [1989], Ludlow proposes that a couple of examples somewhat similar to our current case (in that no more specific comparison class seems to be at work) be handled by supposing "some more general c-class, which is something like objects generally, or perhaps mid-sized earthbound objects" [1989: 530]. But could the case of our director really be handled by such general reference classes? Our director is ready to count as "tall" something that's at least 14 to 16 feet tall. Does that correspond with a good use of "tall for an object generally," "tall for a mid-sized earthbound object," or the like? I, for one, have little idea of how tall an object would have to be to be taller than most mid-sized earthbound objects – or to be taller than the vast majority of such objects, or 80% of such objects, etc. What does the "roster" of those objects and their various heights look like? Before writing this paragraph, I had given no thought to any such questions, and I'm now quite open to supposing that a 14 foot tall object would be taller than the vast majority of the objects in such a class (think of all the objects smaller than that!), and so would satisfy even extremely demanding standards indexed to such classes – but also to the possibility that it would be smaller than most of the objects in such a general class (think of all the buildings, the mountains!), and so would fail to satisfy even very lax standards indexed to such a class. Suppose the director is like me on this matter. She really has no opinions, and would venture no opinion very quickly if asked, how a 14 foot tall object would compare in height with most of (the vast majority of, etc.) the objects in such a general class. It then seems unlikely that by "tall" she meant anything like tall for the objects in such a general class. Indeed, she doesn't seem to be in a position to mean any such thing by "tall." Yet, in our example, that seems to be no bar to her using "tall" in the way we are imagining. It's really looking like we should take her at her word when she says that she means "about 14 to 16 feet tall" by "tall", and suppose her (perfectly fine) use of "tall" is simply governed by a tape-measure standard. Attempts to force such uses into the IRT straightjacket seem implausible.

Likewise, other gradable adjectives can be governed by "tape measure standards" (though the "tape measure" label can't be taken quite as literally). For instance, cases can now easily be constructed by the reader where "rich" doesn't seem to mean anything that has to do with any comparison class (it doesn't mean anything like "rich for an American"), but rather simply means something like "having at least $500,000," or something vaguer than that but which approximates that in its meaning.

I should be clear about an important point of my argumentative methodology used above, but also in what's to follow. Above, I consider what the director would be likely to say if she were asked what she meant by "tall," appealing to the fact that she would give a tape-measure-style answer, rather than cite any comparison class. But I'm not presuming that we must take what speakers say when asked what they mean with inviolable seriousness in accounts of what they mean. What I'm really relying on is that there just is no good candidate for what the relevant comparison class might be. That the speaker wouldn't cite a comparison class to specify what she meant is some helpful corroborating evidence for my conclusion. That she probably couldn't come up with a plausible candidate comparison class if pushed in that direction is better evidence, but in part because it is approaching what I'm really relying on here: that we can't come up with a plausible candidate class, that there just aren't any plausible candidates.[8]

5. Can We Get By Using Only Tape-Measure Standards?

Indeed, once we have these tape-measure standards, one can start to wonder whether we really need the group-indexed ones (at least as a separate category) any more. As I've already noted in the case of statistical group-indexed standards, and as seems to apply to tape-measure standards as well, many of our uses of gradable adjectives seem quite a bit more vague than these numerical clarifications, and when we're pressed on the issue of what we meant, in cases involving both types of standards, we'll often indicate that we meant something vague by combining some device like "something like" with our numerical clarification: "I meant something like taller than about three-quarters of fourth-graders", "I meant something like at

[8] Christopher Kennedy [1999] has argued, on different grounds, for a thesis complementary to my conclusion here. Kennedy concludes that "an empirically and explanatorily adequate semantics of gradable adjectives must introduce abstract representations of measurement – degree qua intervals – into the ontology." Kennedy argues for the need for tape measures (not his term) in our semantics for gradable adjectives because accounts that work with comparison classes but without tape measures cannot account for a variety of comparative constructions. If Kennedy is right and tape measures are involved anyway in the semantics of gradable adjectives, one might then wonder: "Well, what's to stop us from at least sometimes simply using gradable adjectives governed by tape-measure standards, without any comparison class involved? Often, that would be very handy." Well, if "A" quite generally meant "A for an F", the meaning of gradable adjectives is what would stop us from using them with tape-measure standards. But apparently nothing stops us from using tape-measure standards. We do it very often.

least six-foot, two inches tall." Given all this vagueness, one might be tempted by the thought that all our uses of gradable adjectives are governed by (often vague) versions of tape-measure standards, and "for an F" clarifications are just ways of indicating some fairly vague range of tape-measure standards.

But I think this temptation should be resisted. Despite all the vagueness in some of our uses of "tall" (and of other gradable adjectives), we have good reason to think that sometimes we mean "tall" by some (perhaps quite vague, perhaps not) group indexed standard, while on other occasions we mean "tall" by some (perhaps quite vague, perhaps not) tape-measure standard. Consider a man living in Florida, who has a grandson, say, Joey, whom the grandfather hasn't seen since Joey was a baby, because Joey lives far away in Hawaii, and the family is too poor to travel much. Joey's mother calls the grandfather, as is her weekly custom, to get the grandfather caught up on family news. Joey's mother says, "Joey is getting so tall! The doctor says he's in the 98th percentile of 10-year-old boys." The grandfather, we'll suppose, has no idea how tall 10-year-old boys tend to be (it's been so long since the grandfather had sons of his own that age!), no idea how tall one must be to be in the 98th percentile for 10-year-old boys, and no idea, at least in terms of feet-and-inches, of how tall Joey is. Still, Joey's grandfather is now in a position to brag to his golfing buddies about the fact that Joey is "tall." Here, since he has no idea about how tall one must be to be in the 98th percentile for 10-year-old boys, the grandfather doesn't seem to mean even vaguely any tape measure standard use of "tall." He seems to mean the group-indexed "tall for a 10-year-old boy." We'll leave open the issue of <u>how</u> tall for a 10-year-old boy Joey's grandfather use of "tall" means; that may well be extremely vague. He knows that Joey is in the 98th percentile, so he might not have to have very specific intentions regarding what group-indexed standards he means to be using: Since Joey satisfies even extremely demanding group-indexed standards, his grandfather doesn't have to be very careful about what exactly he means. If a cranky buddy is foolish enough to challenge the grandfather's claim – "What'dya mean he's 'tall'?" "I mean he's tall for a 10-year-old boy." "<u>How</u> tall for a 10-year-old boy?" – the grandfather knows he has a killer response – "Well, he's in the 98th percentile!" – that should satisfy the buddy that Joey meets even extremely demanding group-indexed standards for being "tall."

It's natural to think of tape-measure standards as being somehow more "basic" than group-indexed standards, which are in turn naturally thought of as being "up a level" from the

tape-measure standards. For a given group-indexed standard, together with facts about the items in the relevant comparison class, will yield a tape-measure standard, which in some good sense provides a truth-condition for the claim in question. When the group-indexed standard is one involving a certain threshold of the percentage of Fs one must be A-er than to count as A, the only facts about the comparison class that are needed are facts about how tall (or old or rich, etc.) the various members of the class are.[9] So, suppose Joey's grandfather's use of "tall" means some fairly specific group-indexed thing: suppose he means "at least in the 98th percentile for fourth-grade boys." That group-indexed standard, together with facts about how tall the various fourth-grade boys are, will yield a particular quite precise tape-measure standard. And if Joey's grandfather's use of "tall" is instead governed by some quite vague version of a group-indexed standard, that vague standard together with facts about the heights of fourth-grade boys will yield some quite vague tape-measure standard. When the group-indexed standard used is not of the percentile-threshold variety, but rather resembles some of the very different standards that Fara took note of, it should still, together with facts about members of the relevant class, yield some tape-measure standard, but the facts may have to go well beyond just facts about how tall (or old or rich, etc.) the members of the class are.

We have uses for standards of both types (at both levels). For sometimes, as in the case of the director we met earlier in this section, it's very handy to mean something tape-measurey by "tall," since there is no relevant comparison class that immediately presents itself, while on other occasions, as is the case with Joey's grandfather, we're not in a position to mean anything tape-measurey, but are well-positioned to mean something sensible (and sometimes even very precise) by a gradable adjective that is instead governed by a group-indexed standard.

6. Group-Relative Standards

In tape-measure standards, we encountered something down a level from group-indexed standards. There's also a type of standard that often governs gradable adjectives that is up a level from group-indexed standards: what we can call "group-relative" standards. As we've

[9] More generally, this is the case whenever what I'm calling a "roster standard" is used. See note 4, above.

seen, when governed by some group-indexed standard, a use of "X is A" means something like "X is A for an F," where F is some particular comparison class supplied by context. By contrast, when governed by what we will call a group-relative standard, "X is A" means that "X is A for its group," where what context provides in fixing the content is not a particular group or comparison class, but rather, in addition to the fact that a group-relative standard is meant, some way of individuating groups (though this will often be left fairly vague). Using fairly precise standards of each type (since, though artificial, they're easier to work with), then, we can now see that a use of "tall" can express a tape-measure standard like "at least 4′10″," a group-indexed standard like "at least in the 98th percentile for 10-year-old boys", or, now, a group-relative standard like "at least in the 98th percentile for one's age-and-gender group."

It's often quite handy to use group-relative standards – and they seem to be quite unproblematically called "standards." So, for instance, the main campus of the University of Texas might seek to increase diversity of its student population by instituting a program whereby any student graduating from a Texas high school whose SAT test scores place her in the upper 5% of students graduating from her own high school automatically qualifies for admission to the main campus. This is a group-relative standard, and is quite unproblematically called a "standard." (Q: "What standard must a student meet to qualify for admission by this new program?" A: "Her SAT scores must place her in the top 5% of students graduating from her own high school.")

But do such group-relative standards ever govern our use of gradable adjectives? Supposing we use both kinds of standards, it can often be unclear whether a given use of a gradable adjective is governed by a group-indexed or a group-relative standard. Consider Joey's grandfather's use of "tall." Did that mean the group-indexed "tall for a 10-year-old boy" or some group-relative thing like "tall for one's age-and-gender group"? If we ask the grandfather what he meant, he's likely to say he meant "tall for a 10-year-old boy." But maybe he meant the group-relative thing, and was answering by stating how what he meant by "tall" applied to the person he has applying it to. If, instead of an open-ended question about what he meant, we gave the grandfather the two options – "Did you mean 'tall for a 10-year-old boy' or 'tall for one's age-and-gender group'?" – he's likely to not know what to do with that question. Indeed, perhaps it's unclear not only to the grandfather himself, but also to omniscient God, how to answer that question. Maybe if there are separate categories for both group-relative- and group-

indexed-governed uses of gradable adjectives, it's simply unclear, for many of our uses, which category they fall in. If it turned out to be unclear for all of our uses which category they fell in, that might well make us wonder whether we should have two such separate categories. Fans of group-indexed standards might then urge us to get rid of the separate category of uses that are governed by group-relative standards, since there would be no cases that clearly fall in that latter category.

However, if one of the two categories is to be eliminated, it will have to be the ever-popular group-indexed standards that get the ax, for there are cases where it looks like only group-relative standards will do. (Though I don't consider it advisable, the best hope for paring our three categories of types standards governing gradable adjectives is to try cutting the middle category of group-indexed ones, and trying to account for every use with either tape-measure or group-relative standards.)

For, first, there are cases in which speakers quite legitimately use a gradable adjective when they don't seem to be in a position to be using any of the likely candidates of particular group-indexed standards or tape-measure standards, but are in a position to use a suitable group-relative standard. Consider a modification of our story about Joey's grandfather. We've already made him ignorant of how tall Joey is, by tape-measure scales, and of how tall 10-year-olds in general tend to be; all he knows is that Joey is in the 98th percentile for 10-year-old boys. But his ability very sensibly and even in ways fairly precisely to apply "tall" to Joey can survive still more ignorance. Suppose he doesn't even know how old Joey is. We'll keep it in the story that he was truthfully told that Joey was in the 98th percentile for 10-year-old boys, but now suppose he's completely forgotten what age Joey is, and for what age he is in the 98th percentile for height for boys. He still remembers, though, that Joey is in the 98th percentile for boys of Joey's age, whatever that age is. Well, then, he's in a perfectly good position to brag that Joey is "tall." When challenged about what he means, he can explain by appeal to a group-relative standard: "I mean he's tall for boys of his age." And when a cranky buddy makes the mistake of further challenging him ("How tall for boys of his age?"), the grandfather still has a killer response: "Well, he's in the 98th percentile!" Here, the grandfather seems to mean the group-relative "tall for one's age-and-gender group." What group-indexed thing could he instead mean? In his state of ignorance about Joey's age, he's in no position to be saying that Joey is tall for a 10-year-old boy; from the grandfather's point of view, that may well be false. (For all he knows, Joey is 6,

and though he's very tall for a 6-year-old boy, he isn't at all tall for a 10-year old boy.)[10] "But surely he has some idea of how old Joey is!" OK, let's suppose that's so: Suppose that he's certain enough that, say, Joey is somewhere in the range of 5-15 years old. But when he says that Joey is "tall," he can't be meaning that Joey is tall for a boy in that broad range of ages. He's in no position to say that, because, for all he knows, Joey is toward the younger end of that range, say, 6 or 7, and is really quite short for a boy in that broad range, that goes all the way up to 15. "But then he perhaps means that Joey is tall for the lower end of that range – tall for boys aged 5-8, perhaps." Well, anything's possible, I suppose, but Joey's grandfather might be surprised to learn that what he's saying is that Joey is tall relative to a group the grandfather isn't even confident Joey is in.[11] So, while we can't be certain, it really seems that the grandfather means the group-relative reading here.

Second, we get further reasons to recognize that group-relative standards sometimes govern our uses from cases of conjoined applications of an adjective to things in different groups. The reader may add to our story of the grandfather to illustrate this (hint: give Joey a little sister who is very tall for a girl of her age, and have the grandfather brag about both of them in the same sentence following my lead below), but we will change stories here, because Jason Stanley has a nice story that embeds some uses of gradable adjectives that are similar to earlier examples used by Ewan Klein [1980: 15-6] and Peter Ludlow [1989: 520] and that will serve us well here:

[10] It's tempting to think that something must be an F before it can sensibly said that it either is or isn't "A for an F." If so, then that Joey isn't a 10-year-old boy would be by itself, regardless of Joey's height, sufficient to rule out that Joey is tall for a 10-year-old boy. I'm not assuming that here. However odd it may be to say "X is A for an F" when X isn't an F, suppose the statement could still be true. (Suppose, then, that though I'm far more than 10 years old, "He is tall for a 10-year-old" can still be true of me.) Still, from the grandfather's point of view, if the standard is set by 10-year-olds, and Joey may well be just 6, then Joey may well not be "tall" by the relevant standard.

[11] As mentioned in the previous note, there is something at least quite odd about saying "X is A for an F" when A is not an F. Thus, even if a 14-year-old can count as being "tall for a 7-year-old," there would be something quite unusual about applying such a standard to him.

> Suppose John, who is very small for his age, identifies with small things. He has a picture on the wall in his bedroom of an elephant fighting off a much larger elephant. He also has a framed tiny butterfly on his wall. When he is asked why he has both things hung up, he says:
>
> (28) That butterfly is small, and that elephant is small. . . .
>
> Now imagine a picture of a butterfly that's surrounded by much smaller butterflies; it's huge for a butterfly. It's next to a picture of an elephant that's surrounded by much larger elephants. The following is a good description of the situation:
>
> (30) That butterfly is large, but that elephant isn't large. [2004: 134-5]

If we didn't have recourse to group-relative standards, and had to make do with group-indexed standards here, we'd have to say (as Stanley himself concludes) that Stanley's (28) and (30) both involve switching the standards that govern the relevant gradable adjective mid-sentence: The first instance of "small" in (28) means "small for a butterfly" while the second instance is governed by the standard "small for an elephant"; similarly for the two instances of "large" in (30). But (28) and (30) both seem to be perfectly fine sentences to assert in the relevant circumstances, even though there seems to be a rule, or at least a strong presumption, that, where context-sensitive terms are involved, the parameters along which the meanings of such terms vary will be held constant throughout the assertion of a conjunction said "all in one breath."[12] This presumption can be defeated. Unusual or strong stress, intonation or emphasis is often used to signal that the "conversational score" with respect to some parameter is being changed mid-breath. An example of this is how one would say the two 'here's, when saying, while pointing first to one and then to another spot on a roadmap, "The exist we want isn't here, it's here." This same sentence sounds like a contradiction if it is said in a quick, "flat-footed" way, without the unusual "fall-rise" intonation, because if, as our presumption would have it, the content of "here" is kept constant throughout, the sentence does express a contradiction. But when you add the intonation, the presumption is defeated, and the sentence sounds fine. Now, (28) and (30) seem

[12] For some explanation of this presumption, as well as the methodology of using "flat-footed, all-in-one-breath" conjunctions to test for ties (or what Kit Fine called "penumbral connections" [1975: 270]) among potentially context-sensitive terms that the presumption underwrites, see [DeRose 1998: 69-72].

fine in the relevant circumstances even without any use of unusual intonation or emphasis used to signal a change of standards. That gives us reason to think that the "conversational score" isn't being changed in the relevant uses of these conjunctions, and that, in particular, the standards governing the key adjectives in these conjunctions don't change from the first conjunct to the second. But they would have to change if the standards were tape-measure or group-indexed standards. Thus, we have good reason to think that group-relative standards govern the relevant uses of (28) and (30), for if "large" in (30) was governed by the group-relative standard "large for a thing of its kind" (where kinds are individuated – as seems natural to the example – so that butterflies and elephants each count as a kind), it could be used by that very same group-relative standard in both of (30)'s clauses, and (30) could still be true without any switching of standards, for indeed the butterfly in question does, while the elephant in question does not, meet that particular group-relative standard for largeness, for the butterfly indeed is, while the elephant is not, large for its kind.[13]

Stanley himself takes the argument the other way. He uses the correctness of these examples to show that we do often with perfect propriety change standards mid-sentence. But the correctness of the examples shows this only if we don't recognize the possibility of group-relative standards sometimes governing gradable adjectives. Allowing group-relative standards allows us to keep our quite plausible presumption while still making good sense of these examples, since the very same group-relative standard can be governing the two occurrences of the relevant adjective in each of the conjunctions.[14]

---

[13] Stanley's (30), and similar examples, show that "X is A and/but Y is not A" can be true where X is not A-er than Y. Recognizing group-relative standards shows that this can happen where there is no switching of the standard that governs the use of A in the conjunction. "X is A and/but Y is not A" seems to entail "X is A-er than Y" where the two occurrences of A in the conjunction are governed by the same standard that is of the tape-measure or group-indexed variety, but not where a group-relative standard is used.

[14] Stanley also relies on other examples, some of them quite interesting, to show that we appropriately change contents for context-sensitive terms over conjunctions [2004: 134-139]. Addressing these examples one by one would take us too far afield for our current purposes, since the other examples don't involve gradable adjectives. But it is worth looking at one of Stanley's examples here, because holding to our presumption in treating that example can make us recognize a kind of norm that is perhaps important to gradable adjectives as well as to the term, "many", that's featured in Stanley's example. So, consider Stanley's

(32) In Syracuse, there are many serial killers and many unemployed men.

Do we have to suppose that the two instances of "many" are governed by different standards to make sense of this claim? Well, how many xs do there have to be in Syracuse for "there are many xs in

Our case for the conclusion that gradable adjectives can be governed by group-relative standards can be furthered by considering conjunctions that utilize verb-phrase ellipsis, like this variation on Stanley's example:

(28b) That butterfly is small, and that elephant is, too.

Those who don't feel much pressure to accept that there's a presumption generally against changing standards between two uses of an adjective within one sentence may feel a bit more pressure to suppose one can't appropriately vary standards from one clause to the next over verb-phrase ellipsis.[15]

But even (28b), one might think, has two occurrences of "small" – though the second one is implicit – and why couldn't they be governed by different standards? So, third, perhaps a

---

Syracuse" to be true? If you're looking for a number, there seem to be many numbers that seem about as likely as any others to be the threshold that would make (32) true when that number sets the threshold for both uses of 'many'. Yet there's some plausibility in Stanley's suggestion that (32) can be used in such a way that it takes more unemployed men to make its second part true than it takes serial killers to make its first part true. (I'd find the suggestion still more plausible if I could get myself to hear In Syracuse, there are many serial killers, but not many unemployed men as being potentially OK (appropriate/correct, as well as true) despite the fact that are more unemployed men than serial killers there. But I can't.) But when I suppose that Stanley's right, that doesn't make me think that there is a change in the standards governing 'many' from one use to the other in (32), but rather inclines me to conclude that 'many' can be governed by standards that aren't simply a matter of the number of things in question – perhaps something like the "norms of expectation" that Fara writes about [2000: 67], so that both of the 'many's mean something like 'surprisingly many.'

[15] Ludlow [1989: 520-4] and Stanley [2000] attempt to handle examples like (28), (30), and (28b) in another way. Stanley very compactly summarizes their approach at [2003: 272, fn. 5]. Ludlow and Stanley assume in their treatments that uses of gradable adjectives have a variable for comparison classes in their logical form, and in section 5, above, we've already seen reason to doubt that. At least insofar as I have a good feel for what's meant by "logical form" in such discussions, the very "logical form" of "X is A" cannot have a variable for comparison classes if the meaning of such claims involves comparison classes only on some resolutions of its context-sensitivity, and there are good uses of it where there simply is no relevant comparison class that is a component of the meaning.

I suspect that if my pluralism is correct about gradable adjectives, the role that "logical form" is often given in explaining various "binding" phenomena will have to be revised. The binding phenomena appealed to by the likes of Ludlow and Stanley are real; I don't question the existence of readings they claim (at least in most cases). But I think we have good reason to suppose that often the variables being bound are not present in the syntax or logical form of the claims in question, but get introduced only as context-sensitive elements in the sentences take on particular values. (Where such a bound reading of the sentence is intended, that makes salient a reading of the context-sensitive element that provides the variable to be bound.) But I also think we have other reasons not dependent on pluralism regarding gradable adjectives for drawing the same conclusion, anyway.

better case for the necessity of recognizing the use of group-relative standards comes from collective uses of gradable adjectives. So, suppose John has a scrapbook of pictures of small things. It contains pictures of the unusually small elephant, the tiny butterfly, his family's old Volkswagon Beatle (very small for a car), the small house his family used to live in, etc. The reason he has pictures of these things in his scrapbook is that each is small relative to other things of its own kind. When asked what all the things depicted in his scrapbook have in common, John can quite appropriately respond:

(28c) They are small.

Here we have one occurrence of "small," but John's claim can only be true – which it surely is – if all the things depicted in the scrapbook meet the standard that governs this occurrence. But surely that standard is here the group-relative "small for a thing of its kind." All the items meet that standard, and that they meet that standard is the reason John put them in his scrapbook. There are of course some tape-measure and group-indexed standards that all the items meet as well, but none of those standards seem to be relevant here.[16]

To get an example from this story of a use of "small" applied to an individual thing (something of the form "X is A") that is pretty clearly governed by a group-relative standard, suppose you now hear that John has added a picture of a new item to his collection. You might not know what kind of thing is being added to the collection, and so may not be in a position to use any plausible group-indexed (or tape-measure) standard to describe its size, but based just on your knowledge of what kind of thing makes it into John's collection, you seem in a position to quite sensibly assert

(28d) The thing John just added to his collection is small,

using the group-relative "small for a thing of its kind" as your standard for "small," and we have an example in some key ways structurally similar to the example involving the grandfather's use of "tall" with which we began this section.

[16] If a die-hard defender of IRT were to insist that "small" here is indexed to the smallest type represented in the book – say, butterflies – we should observe that if we were to ask John whether an ant which is not as large as his butterfly though it is quite large for an ant is "small" and a picture of it belongs in his scrapbook of "small" things, he'll say no.

7. Mixed Standards

Finally, it's worth pointing out that there seem to be perfectly fine uses of gradable adjectives that are governed by standards that combine elements of different of the types of standards we have looked at, or so it seems to me. For instance, when the coaches of an NBA team say they're looking to draft "tall" players this year to add to their team, they can mean something like "tall for one's position, and at least 6′8″," which combines a group-relative with a tape-measure standard. A 6′6″ point guard, even though he is tall for his position (even by NBA standards) doesn't count as "tall" because he is not at least 6′8″. Nor does a 6′10″ center make the cut, since he is not tall for his position. But a 6′9″ point guard does count as "tall". (The next Magic Johnson, perhaps!)

Such mixed standards are often very useful for certain purposes. The above mixed standards for "tall", for example, can be useful for the coaches of a team badly in need of help in rebounding. They can be aided by players who are tall for their position, but no player that falls below a certain tape-measure height can be much help, even if he is tall for his position. One way for speakers to get into contexts that are at least roughly governed by such standards, other than by explicitly stating that by "tall" one will mean "tall for one's position, and at least 6′8″," is by having various proposals shot down and others accepted:

"C'mon! He's not tall. Alright, so he's a point guard, but he's only 6′6″. How's that gonna help us on the boards?"

"What? A 6′10″ center? That's only average for an NBA center. We need <u>tall</u> players. C'mon!"

"A 6′9″ point-guard? Now you're talking! <u>That's</u> the tall player that can help us."

8. Pluralism and Imprecision

Where does all this leave us? According to IRT, where "A" is a gradable adjective, "X is A" quite generally means that "X is A for an F," where the comparison class, F, can vary with and is supplied by context. But we've seen that IRT, even in its weak form, is wrong. "X is A" can mean something like "X is A for an F," but that's only one of the kinds of thing it can mean. In addition to group-indexed standards, there are uses of gradable adjectives that are governed by quite different standards – standards of the group-relative and tape-measure varieties. (And even where "X is A" does mean "X is A for an F," context's work is not necessarily done, for "X is A for an F" can mean seriously different things in different contexts.)

It seems we should replace IRT, or the "for an F" myth, with a pluralist approach to gradable adjectives, according to which these terms can be governed by a variety of different kinds of standards. Because of this variety, with the group-indexed ones being just one type among several, the most we can give as an account of what "X is A" generally means is the minimalist account that it means that "X is A by standard S."[17]

The way forward from that minimalist account would seem to be to discern what types of standards can govern the use of gradable adjectives, and how these different standards work. I've here tried to identify some of the different types of standards that are used, but I realize I've made scarcely a beginning on the task of trying to figure this all out. But what we've seen already may help to clear the way a bit for contextualism about knowledge attributions to get a fairer hearing than it would if the "for an F" myth continued to be widely assumed among philosophers casually interested in gradable adjectives. For, first, the two worries addressed in section 4 are assuaged already by the demise of the strong version of IRT, for the reasons given in that section.

And second, if the most we can give as an account of what "X is A," where A is a gradable adjective, generally means is the minimalist account that it means that "X is A by standard S," then what contextualists in epistemology propose no longer looks so foreign, for a quite similar minimalist, standards-based account is available for knowledge attributions, too: "Y

[17] Though it's natural to use "minimalist" to describe this proposal, note that this "minimalist" account is very different from the thesis of "semantic minimalism" championed recently by Herman Cappelen and Ernie Lepore (see esp. [2005]). Cappelen and Lepore's "minimalism" explicitly denies the very context-dependence in gradable adjectives that I seek to understand.

knows that P" means that "Y knows that P by standard S." This isn't a great similarity: All it amounts to is that each is governed by varying standards. What's more, one will find nothing like analogous of group-indexed and group-relative standards involved in the various uses of knowledge attributions; all the standards that govern our knowledge-talk seem to be more closely analogous to what we've been calling tape-measure standards. But I don't mean to be basing a positive reason for accepting contextualism for "knows" on this similarity, or even to be laying the groundwork for such a case. But, to the extent that assumed but mistaken theories about gradable adjectives have been a hindrance here, I do hope to have removed an obstacle to accepting contextualism about "knows" and other contextualisms as well.

REFERENCES

Cappelen, Herman and Lepore, Ernie 2005. Insensitive Semantics, Basil Blackwell, 2005.

DeRose, Keith 1998. Simple *Might*s, Indicative Possibilities, and the Open Future, The Philosophical Quarterly 48: 67-82.

Fara, Delia Graff 2000. Shifting Sands: An Interest-Relative Theory of Vagueness, Philosophical Topics 28: 45-81.

Fine, Kit 1975. Vagueness, Truth and Logic, Synthese 30: 265-300.

Hawthorne, John 2004. Knowledge and Lotteries, Oxford: Oxford University Press.

Kennedy, Christopher 1999. Gradable Adjectives Denote Measure Functions, Not Partial Functions, Studies in the Linguistic Sciences 29: 65-80.

Klein, Ewan 1980. A Semantics for Positive and Comparative Adjectives, Linguistics and Philosophy 4: 1-45.

Ludlow, Peter 1989. Implicit Comparison Classes, Linguistics and Philosophy 12: 519-533

Ludlow, Peter 2005. Contextualism and the New Linguistic Turn in Epistemology, in Contextualism in Philosophy: Knowledge, Meaning and Truth, ed. Gerhard Preyer and Georg Peter, Oxford: Oxford University Press: 11-50.

Richard, Mark 1995. Defective Contexts, Accommodation, and Normalization, Canadian Journal of Philosophy 25: 551-570.

Schiffer, Stephen 1995. Descriptions, Indexicals, and Belief Reports: Some Dilemmas (But Not the Ones You Expect), Mind 104: 107-131

Stanley, Jason 2000. Context and Logical Form, Linguistics and Philosophy 23: 391-434.

Stanley, Jason 2003. Context, Interest-Relaitivity and the Sorites, Analysis 63: 269-280.

Stanley, Jason 2004. On the Linguistic Basis for Contextualism, Philosophical Studies 119: 119-146.