

Single Scoreboard Semantics

Keith DeRose

Yale University

This paper concerns the general question of what happens to the “conversational score” in cases where the participants in a conversation are pushing the score in different directions. Given my own epistemological interests, I am particularly interested in the special case of knowledge attributions, and since our question is best addressed by means of particular examples anyway, I’ll for the most part frame the discussion in terms of that special case.

So, suppose a skeptic presents a skeptical argument to the conclusion that her conversational partner does not know something that she would ordinarily take herself to know, and that, in so doing, at least according to typical contextualist analysis, the skeptic executes maneuvers which have a tendency to inflate the standards for knowledge to a level at which her partner indeed does not count as knowing the thing in question. Suppose, however, that our skeptic meets with an “Aw, come on!” response from her opponent, who insists, and continues to insist, that she does indeed know the thing in question. According to the contextualist, who wins? Who is speaking the truth? The skeptic, her opponent, both, neither? What happens to the truth conditions of knowledge attributing and denying claims in cases where the parties to the discussion seem to be pushing the “conversational score” in different directions?

In previous work, I have remained neutral about such questions. Here, I’ll present the answer I now favor and my reasons for favoring it, after laying out several other options, a couple of which are at least suggested by David Lewis in his very influential “Scorekeeping in a Language Game,” which, of course, inspired my use of “scores” and “scoreboards” in thinking of these matters.¹ I’m very interested in what others think about this question: How other contextualist are inclined to answer the question, and what

non-contextualists think is the contextualist's best answer. I hope that laying out the issue and some of the options will encourage some good thinking on this topic.

The Type of Situation Addressed in this Paper

Many seem to assume that the contextualist's answer will be that both the skeptic and her opponent are speaking the truth. The skeptic's claim that her opponent "does not know" is true iff her opponent fails to meet the extraordinarily high standards that the skeptic's claim has a tendency to put into place. Since her opponent does not meet those standards, the skeptic's claim is true. Her opponent's insistence that she "does know" is true iff she meets lower, ordinary standards for knowledge. Since she does meet those lower standards, her claim is true, too. Both are speaking the truth, and they are failing to contradict one another. They are talking past one another.

There are certain ways the debate could go which I think would make that the correct answer to our question. For example, suppose the skeptic says, "You don't know, and by this I mean...", and completes the sentence by explaining the very high standards her opponent would have to meet before the skeptic would count her as a knower. Her opponent replies, "I do know, and by this I mean...", completing the sentence by specifying the moderate epistemic standards that she is claiming to meet. Both speakers having specified what they mean, and having specified very different meanings, they stop thinking of themselves as contradicting one another. Even if I were an invariantist, under these circumstances I would think that both the skeptic and her opponent are speaking the truth. As an invariantist, I'd think that the truth-conditions of standard uses of "know" do not vary with context. But when a speaker resorts to an explicit use of a "and by this I mean..."-like construction, I'm inclined to think that the speaker's utterances come to mean what the speaker says it means. Speakers are free to stipulate what they will mean by a new term they are introducing, and they're free to stipulate a special meaning they are giving to an existing term. Thus, if I were an invariantist, I'd think that in the imagined situation, at least one of our two speakers is using "know" in a special, non-standard way, and that both the skeptic and her opponent are speaking the truth. (But I

would also think that such a conversation shows nothing about standard uses of “knows”.)

Different situations that do not include anything that looks like a stipulation of (perhaps non-standard) meaning also tempt me, as a contextualist, toward the same verdict. Suppose, for instance, as often happens in such debates, that one or the other speaker, or some third party, advances some contextualist, or proto-contextualist, analysis of the debate that is transpiring, phrased in terms of what each party of the debate “means” by “knows,” saying something along the lines of, “Well, by ‘knows’, you seem to mean..., while I seem to mean...”. Suppose those involved in the debate all accept this analysis, and our speakers, now believing that they mean different things by “knows,” consequently stop thinking that they are contradicting one another. I won’t venture a guess as to what invariantists will say about a case like this, where the participants to a discussion have explicitly accepted such an analysis of their own debate, but have not said anything that looks like an explicit stipulation of meaning. But as a contextualist, as I’ve admitted, I’m at least tempted to conclude that both the skeptic and her opponent are speaking the truth in such cases. But that is not the kind of case I’ll be addressing in this paper.

It often happens in debates between skeptics and their opponents that nobody even offers such an analysis of what’s transpiring, much less is such an analysis accepted; the skeptic and her opponent do take themselves to be contradicting one another; each intends to be contradicting what the other is saying; and, beyond what’s going on privately in their own minds, each is publicly indicating that they are (or at least mean to be) contradicting the other, by saying such things as, “No, you’re wrong. I do know.” It’s such cases as this that I’m addressing in this paper. Perhaps some of us here have been party to such debates. And about this kind of case, I am not inclined to think that both our speakers are speaking the truth, but failing to contradict one another.

The type of case in question is one in which there is a conversation crucially involving a context-sensitive term (one that can express different specific contents, given context) and in which, as we can put it, the personally indicated content of one speaker – the content that speaker’s conversational maneuvers have a tendency to put into place for that term – diverges from the personally indicated content of the other speaker, but in

which the speakers still indicate that they are contradicting one another.² This is a type of situation that can arise with other context-sensitive terms, too, and so, though I'm here addressing what to say about such situations where "knows" is involved, the question of what a contextualist should say about such a case is a quite general one, and so one quite appropriate to a conference on "contextualism in epistemology and beyond."

I should note that I won't be getting into any specifics about the precise mechanisms by which speakers can change the score or resist such changes. With regard to "knows", different contextualists have proposed different accounts of the rules by which the standards for knowledge can be raised, and, though we haven't been specific about the precise means, we all seem to suppose that there are ways that standards can be lowered or raises in epistemic standards can be resisted. We will just suppose our skeptic has executed a maneuver (whatever that maneuver is and however exactly it works) that has a tendency to raise the epistemic standards, and that her opponent has responded by executing a maneuver that has at least some tendency to keep lower, ordinary standards in place. We are then concerned with the question of what happens to the conversational score in such a situation. For simplicity, I will (except for some brief speculation in note 9) concentrate only on cases in which there are just two speakers involved in a conversation. Also, I will ignore complications caused by vagueness in the personally indicated content of certain speakers, and suppose that each speaker is personally indicating a precise set of standards.

Multiple, Personal Scoreboards

Many seem to think or assume that, even with respect to the cases just described, the contextualist's answer, or the contextualist's best answer, to our question will still be that both of our debaters are speaking the truth, and they are failing to contradict one another. The idea here is that the truth-conditions of each speaker's spoken claims will directly match the personally indicated content of that speaker. Thus, since the skeptic is executing conversational maneuvers that tend to put into place the high, skeptical standards for knowledge, her claims that her opponent "doesn't know" are true iff her opponent fails to meet those extraordinarily high standards. And since her opponent's conversational maneuvers tend to put (or keep) in place lower, more ordinary standards,

her claims that she does indeed “know” are true iff she meets those lower standards. The picture seems to be that each speaker, in addition to having their own personally indicated content (having certain standards that their conversational maneuvers have at least some tendency to put in place) also have their own personal scoreboard, by which I mean that the truth conditions of each speaker’s use of “knows” is particular to that speaker, and presumably matches that speaker’s personally indicated content.³ Of course, it helps various communicative purposes if speakers engaged in a conversation adjust to one another’s usage, and come to have matching scores on their scoreboards. And those who think in terms of multiple, personal scoreboards will probably think that’s what usually happens. But in the cases we’re considering here, that doesn’t happen; our speakers fail to adjust to one another. And if you’re a contextualist who is thinking in terms of each speaker having their own personal scoreboard, you will conclude that both our skeptic and her opponent are speaking the truth, and are failing to contradict one another.

This position, of course, does have its appeal. There is something to be said for thinking the truth conditions of a speaker’s knowledge claims match the epistemic standards that she herself is indicating by her conversational maneuvers.

Single Scoreboard Semantics

But there is a cost, too. Each speaker, in addition to indicating certain epistemic standards, also indicates that they are contradicting the other speaker. And the multiple scoreboards position has the unfortunate result that our two speakers are not, as they surely seem to be, and as they take themselves to be, contradicting one another.

One thought that can lead one to the multiple scoreboards view is that we should respect what each speaker is indicating in assigning content to the claims of that speaker; speakers should be in control of their own meaning. But, as I’ve just noted, in the cases in question, each speaker is indicating two different things: that the standards be such-and-such, and that their claims contradict those of the other speaker. But where the indicated standards of the two speakers diverge, we cannot consistently respect all of the indications being given by our speakers. Why should it be the clear indications being

given by each speaker that her claims be understood to contradict those of the other that give way here?⁴

My own thought about the cases under consideration has always assumed what I call “single scoreboard semantics,” due to the influence of David Lewis’s “Scorekeeping in a Language Game,” which seems to me to promote such an account.⁵ On this view, there is a single scoreboard in a given conversation; the truth-conditional content of both our speakers’ uses of “knows” are given by the score registered on this single scoreboard. The score of course can change as the conversation progresses, and the score it registers is responsive to the maneuvers made by all the various speakers, but there is at any given time a single score that governs the truth conditions of all the speakers’ uses of the relevant term. Thus, both our skeptic’s and her opponent’s use of “knows” are governed in their truth conditions by what that single scoreboard registers when their claims are made. This picture, of course, promotes the thought that our skeptic and her opponent really are contradicting one another: If there is a single scoreboard, then our skeptic is denying precisely what her opponent is affirming as they debate back and forth.

Of course, if we imagine the score moving sharply up and down throughout the conversation – moving up (the standards for knowledge going up) every time the skeptic speaks, and dipping down suddenly whenever her opponent makes a claim – the result would be equivalent to supposing that each speaker had their own personal scoreboard. So that’s not the idea. There are tricky questions we’ll see later about the exact timing of certain changes of score which can make it difficult for the contextualist who embraces single scoreboard semantics to know what to say about the truth conditions of some of the opening claims made in our debate. But once the skeptic has made her standards-raising maneuvers a time or two, and her opponent has responded with her stubborn, resistant maneuvers, and they continue to debate, “You don’t know”, “I do know”, the conversational score has presumably reached whatever state of equilibrium it reaches in such situations, and, at least as far as the truth-conditions of their knowledge claims go, on the single scoreboard picture, the skeptic is denying precisely what her opponent is claiming. So they are contradicting one another.

But which is speaking the truth? What is the conversational score in such a situation? Let’s canvass some possible accounts of what happens to the conversational

score – what happens to the truth conditions of claims about what is and is not “known” – in the situations we’re considering.

The Skeptic Wins

In my experience, the view that is most often ascribed to contextualists – because it is thought this is what we do think, and/or because it is thought that this would be the best view available to us – is “multiple scoreboards” view that both the skeptic and her opponent are speaking the truth. As I’ve already noted, I do not in fact accept this view. I’ve also asked Stewart Cohen, who has informed me that he doesn’t accept that answer either. Nor do either of us accept the view that is the next most commonly ascribed to the contextualist: that, in the situations under consideration, the skeptic’s extraordinarily high standards prevail, so that the skeptic is speaking the truth when she charges, “You don’t know,” and her opponent is saying something false when she insists, “I do know.”

Some of what David Lewis wrote suggests this very skeptic-friendly view. Background: Lewis worked within the relevant alternatives picture of knowledge, according to which an attribution of knowledge of a proposition *p* to a subject *S* requires for its truth that *S* can rule out or eliminate all the relevant alternatives to *p*. The skeptic changes the score by enlarging the range of alternatives to *p* that are relevant so that they include various hard-to-eliminate alternatives that are usually irrelevant. In “Scorekeeping in a Language Game,” Lewis writes:

The commonsensical epistemologist says: “I know the cat is in the carton – there he is before my eyes – I just can’t be wrong about that!” The skeptic replies: “You might be the victim of a deceiving demon.” Thereby he brings into consideration possibilities hitherto ignored, else what he says would be false. The boundary shifts outward so that what he says is true. Once the boundary is shifted, the commonsensical epistemologist must concede defeat. And yet he was not in any way wrong when he laid claim to infallible knowledge. What he said was true with respect to the score as it then was.⁶

It isn't completely clear what Lewis meant by saying that the commonsensical epistemologist "must concede defeat." What if, as in the situations we are here considering, the epistemologist doesn't do what he "must" do, and stubbornly repeats his claim to know? It's tempting to (and I'm inclined to) read Lewis here as holding that the commonsensical epistemologist would then be saying something false, because, in this context, I am tentatively inclined to understand "the commonsensical epistemologist must concede defeat" to mean that, if he says either that he "knows" or that he doesn't "know," he must choose (admit) the latter if he is to speak the truth.⁷

In his later paper, "Elusive Knowledge,"⁸ Lewis seems to accept an analysis that also yields that result. In that paper, the irrelevant alternatives are the ones that are "properly ignored"; the relevant alternatives that must be "eliminated" by one's evidence if one is to count as a knower, then, are those that are either not ignored, or are ignored but improperly so. The skeptic, by calling attention to various skeptical hypotheses, creates a situation in which those uneliminated possibilities are not ignored at all, and so are not properly ignored.⁹ If the skeptic's opponent were to claim to know, then, it seems, he would be saying something false.

Perhaps Lewis's apparent backing of the verdict that the skeptic wins is at least largely responsible for its popularity. Another possible contributing cause is that other contextualists, as well as Lewis, have concentrated our efforts on discerning the mechanisms by which epistemic standards are raised, and have not worried that much (at least in print) about how standards can be lowered, or threatened raises in standards can be resisted.

Does It Matter if the Skeptic "Wins"?

One common complaint against the contextualist approach to skepticism is that it is too friendly to the skeptic. To some extent, that complaint is unavoidable. For just about any contextualist will hold that we don't know according to the epistemic standards that the skeptic's maneuvers have at least some tendency to put in place, and that they actually succeed in putting in place when their moves are not resisted, and that by itself will be too soft on skepticism for the tastes of many. But the views we have considered so far make the situation worse for the skeptiphobic. For on the multiple scoreboards view, the

skeptic truthfully states her conclusion even when she is resisted – though her opponent is also held to be speaking truths when she insists that she does “know”. And on the only single scoreboard view that we have looked at so far, things are even worse: Not only is the skeptic speaking the truth, but her opponent is speaking falsely.

That the two skeptic-friendly views we have looked at to this point are so often assumed to be what the contextualist will say has probably made contextualism unattractive to many, and it might help make the contextualist case, at least to many, to point out that there are other, less skeptic-friendly, ways to take the contextualist approach.

But before considering these other options, including the one that I favor, I’d like to register my view that, though I think the skeptic does not “win” in the situations we are considering – I don’t think the skeptic is speaking the truth when she says we do not “know” in situations in which she is resisted by a stubborn opponent –, it would not matter that much if she did win. Once we see that we don’t know according to the standards the skeptic’s maneuver have at least some tendency to put in place, but that we do know according to the ordinary standards that govern most of our thinking and speaking about what is and is not “known,” so that the skeptic’s “success” has no tendency to show that we’re usually deeply mistaken in our ordinary thought about what we do or do not “know,” we have seen most of what’s important in the contextualist approach. We can then go on to discuss the somewhat important matter of how serious or worrisome a result it is that we don’t meet the skeptic’s extraordinarily high epistemic standards, once we’re freed from the concern that that shows our ordinary thought about knowledge to be mistaken. (My own view here is that it is not this is not very worrisome at all.) But, once all the above is seen, I myself don’t find it all that important who “wins” – who ends up speaking the truth – in various debates between skeptics and their opponents in which the parties do not see the contextualist resolution to their debate, and fail to adjust to one another’s standards.

But those who see things differently, and are especially bothered by the thought the skeptic “wins” debates that go the way we’re considering here, may be especially interested in some of the less skeptic-friendly options we’ll look at below. (For those who, like me, don’t much care whether the skeptic “wins” in the sense in question, this

can still be an interesting question about how to understand conversations involving divergence in personally indicated content in context-sensitive terms.) We will start with a view that is very unfriendly to the skeptic.

Veto Power

As I've noted, Lewis's contextualist writings tend toward the view that, in the situations we are considering, the skeptic wins. But in a section of "Scorekeeping" different from the one we've looked at already, Lewis at least suggests a different view – not about knowledge, but about "standards of precision":

Taking standards of precision as a component of conversational score, we once more find a rule of accommodation at work. One way to change the standards is to say something that would be unacceptable if the standards remained unchanged. If you say "Italy is boot-shaped" and get away with it, low standards are required and the standards fall if need be; thereafter, "France is hexagonal" is true enough. (p. 352, emphasis added)

Important here is the phrase I've emphasized. What if our speaker does not get away with lowering of the standards of precision? Then, it at least seems to be suggested, the standards do not fall. Our speaker has said something that – in this case, by a "rule of accommodation" – has a tendency to lower the standards of precision so that Italy counts as "boot-shaped." The picture at least suggested, though, is that those he is talking with have "veto power" over this changing of the conversational score: If they don't let him get away with changing the score, then he does not succeed in changing the score.

A similar position is tempting about our debate between the skeptic and her opponent. According to the contextualist, the skeptic has executed maneuvers that have a tendency to raise the standards for knowledge. But in the situation we are considering, her opponent does not accede to this raising of epistemic standards; the skeptic is not allowed to "get away with" changing the score. On the view under consideration, then, the standards are not raised. Here, we get a view on which the skeptic's opponent ends up speaking the truth, and the skeptic's claims are false. This depends on the assumption

that the epistemic standards start out meetably low, perhaps because they have such a “default” value, but that is a plausible enough supposition.

There is a tricky issue about timing here, that I’m not sure what to say about. Suppose our skeptic has just made her standards-raising maneuver, but her opponent has not yet responded. Her opponent is quickly deciding between two responses, one of which would constitute letting the skeptic get away with raising the standards, the other of which would constitute not letting her get away with it. But right now, in the brief moment between before the skeptic’s opponent speaks, where are the standards? Was the skeptic’s knowledge-denying utterance true? Can the truth value of that claim depend on something – how her opponent will respond – that has not yet occurred? I will here remain open about such questions – and similar questions that can arise for other views we’ll be looking at. But I will assume that at least after our skeptic and her opponent have each had a couple of turns, the conversational score has reached whatever state of equilibrium it reaches in such cases of non-cooperation. According to view currently under discussion, since the skeptic has failed to get away with raising the epistemic standards, and has therefore failed to raise them, it is her opponent is speaking the truth when she insists she does “know”, while the skeptic is speaking a falsehood when she claims her opponent does not “know”.

Reasonableness Views: Pure Reasonableness and Binding Arbitration

To some, the matter of what would be a reasonable score, given the situation of a conversation, is important in determining what the score actually is in cases like those we are considering. The simplest way for this to work is what we can call the “pure reasonableness” view, on which the conversational score – the relevant aspect of the truth conditions of the various speakers’ claims – is whatever would be the most reasonable score for the speakers to use, no matter what the personally indicated content of the various speakers may actually be.

Another way reasonableness can figure in is on what we may call the “binding arbitration” model. In major league baseball, at least as I understand the process, when a player and his team go to binding arbitration to decide the player’s salary, each makes a bid or proposal on what the player’s salary should be, the player of course, suggesting a

higher figure than the team. Evaluators then decide which of the two bids is the more reasonable, given the abilities of the player and the salaries of other players at the same position, and other factors, and that becomes the player's salary. On the "binding arbitration" model of conversational scorekeeping, when the personally indicated content of one speaker diverges from that of the other that she is talking with, then the score – which gives the truth conditions of both speakers' claims – matches the personally indicated content of the speaker who is indicating the more reasonable content of the two.

The binding arbitration view has some attractions. On it, as on other single scoreboard views, our skeptic and her opponent really are, as they seem to be, contradicting one another. And the personally indicated content of one speaker – the one that is indicating the more reasonable standards – matches the truth conditions of her own claims. Of course, the truth conditions of the other speaker's claims don't match up with her own personally indicated content, but that can seem like just punishment for her putting in the more unreasonable "bid."

On both of the reasonableness views we've considered, the matter of which of our speakers is speaking the truth depends on facts about what are the reasonable standards for them to use in their situation. On the binding arbitration model, whichever of our speakers has the more reasonable personally indicated standards of the two is speaking the truth. On the pure reasonableness view, it depends on whether the skeptic's opponent meets the most reasonable standards that our speakers could use given their situation. And, of course, this opens up a whole host of questions that I won't even begin to address about what makes standards the reasonable ones to use. But anti-skeptics who think that the skeptic's standards for knowledge are entirely unreasonable and that we do meet any reasonable standards for knowledge will see in these views a way to deprive the skeptic of victory.

For my part, even if I could see past the problems involved in deciding what are reasonable standards, I wouldn't be attracted by reasonableness views of the type we have been considering. It's good for speakers to use reasonable standards (or, more generally, reasonable scores), of course. But if they opt for unreasonable standards, I'm inclined to think the truth conditions of their claims then reflect those unreasonable standards that they are indicating. In uncooperative conversations where the different

speakers are personally indicating different standards, if one of the speaker's indicated standards can be seen to be more reasonable than the other's, then the more reasonable speaker is in some way more conversationally praiseworthy than the other, I suppose, but I don't see that this is so in a sense that would imply that the truth conditions of both speaker's claim should be thought to take on the truth conditions that match her indicated standards.

The Exploding Scoreboard

We have not yet arrived at the view I prefer. To prepare for it, we'll first consider another view, that, having grown up as a sports-crazy kid on the South side of Chicago during the days when Bill Veeck owned the White Sox and Comiskey Park, I can't resist calling the "exploding scoreboard" view. On this view, when the personally indicated content of two speakers in a single conversation diverge – or at least when they diverge by as much as happens in our debate with the skeptic and her opponent – the scoreboard explodes: There is no correct score, and claims involving the relevant term are neither true nor false.

We are, after all, considering a conversation that is quite defective. Speakers engaged in conversation should adjust to one another's score so they are meaning the same thing by the key terms in question. When this doesn't happen, and especially when the divergence in personally indicated content is great, as in the situations we are here considering, perhaps semantic hell breaks loose.

On this view, neither the skeptic nor her opponent is speaking truths in the situations we are considering – nor is either speaking falsehoods. I'll let you decide for yourself whether and to what extent the skeptic "wins" on this view – and whether that's a problem. She doesn't truthfully claim that we don't "know," but her destructive purposes are served by her creating a context in which her opponent cannot truthfully claim that she does "know".

The exploding scoreboard view, initially at least, strikes me as very plausible when it's applied to cases of great divergence. It's far less plausible when applied to cases of slight divergence.

Consider an example involving the matter of how far out “here” reaches in a given case. You and a colleague are in a giant room – say the grand ballroom of a hotel at the APA “smoker” – and are discussing whether or not Frank is “here.” By “here”, neither of you means to be designating an area so small that Frank would count as “here” only if he were sitting right at the same table that you and your colleague are sitting at. (It’s obvious that he’s not right at your table, so there would be no point in wondering, as you are, whether Frank is “here” if that’s what “here” designated.) But neither do either of you mean to designate an area so large that the whole ballroom counts as being “here.” (You both assume that Frank is somewhere in the ballroom, and so wouldn’t be questioning whether Frank is “here” if “here” included the whole room.) You both mean to designate – and make conversational maneuvers indicative of this – something in between those two extremes. But the area that you persistently indicate is slightly larger than the area your colleague indicates, and you fail to adjust to one another’s usage. (Don’t ask me how this – one of you indicating just a slightly larger area than the other – could happen. Just suppose it does.) While this conversation is slightly defective (the two of you should somehow adjust to another’s usage so that you mean the same area by “here”), it’s not all that bad, and if Frank is in fact in the far corner of ballroom, outside of what either of you have been indicating as the range of “here”, then it seems very plausible that in your conversation with your colleague, “Frank is here” is false, and “Frank is not here” is true. Thus, it is very implausible to suppose that, because of the slight divergence in personally indicated content in your conversation, the “here” scoreboard has simply exploded.

The “Gap” View

The single scoreboard view that cases of small divergence strongly suggest to me is the “gap” view, according to which, as applied to your conversation about whether Frank is “here”:

-claims that “Frank is here” are true (and claims that “Frank is not here” are false) iff Frank is in the region that counts as “here” according to both speaker’s personally indicated content;

-claims that “Frank is not here” are true (and claims that “Frank is here” are false) iff Frank is in the region that does not count as “here” according to either speaker’s personally indicated content; and

-claims that “Frank is here” and that “Frank is not here” are neither true nor false if Frank is in the region that counts as “here” according to one speaker’s personally indicated content, but does not count according to the other’s. (In this case, in a fairly literal sense, Frank is neither “here” nor “there”.)

Applied to cases involving “knows,” then, in cases of small divergence in personally indicated epistemic standards, “S knows that P” is true (and “S doesn’t know that P” is false) where S meets the personally indicated standards of both speakers; “S doesn’t know that P” is true (and “S knows that P” is false) where S fails to meet either set of standards; and where S meets one set of standards but fails to meet the other, both “S knows that P” and “S doesn’t know that P” go truth-value-less.

There’s a lot to be said for giving a uniform treatment to cases of small and large divergence – especially since it will be difficult to draw a line between the two. So, having accepted the gap view for cases of slight divergence, I’m inclined to accept it as well for cases of great divergence. And why not? Even where the divergence is great, isn’t it plausible to suppose that “S knows that P” is false if S fails to meet the personally indicated standards of either speaker? Just so, even in our debate between the skeptic and her opponent, it seems plausible that either speaker would be saying a truth if they were to claim that “Frank does not know that the Cubs won the 1908 World Series” if Frank fails to meet even the standards for knowledge of the skeptic’s opponent.

Indeed, the gap view is the view I’m most inclined to accept about the situations we are here addressing. Its main attraction to me is its impressive ability to simultaneously respect both the sense that our two speakers are contradicting one another

and the feeling that the truth conditions of each speaker's assertions should match that speaker's personally indicated content.

The view does an excellent job of respecting the sense of contradiction – as well as can be hoped for in cases involving gappy claims. Suppose the person you're talking with makes a statement with "gappy" truth-conditions. Take, for example, a standard case of vagueness. She says "Frank is tall," where her claim is true if Frank's height is in a certain range, is false where Frank's height is in a certain lower range, and, we'll suppose, is neither true nor false if Frank's height is in an intermediate range. You reply, testily, "Frank is not tall!" What content can we assign your claim that would do the best job possible of making your claim contradict your friend's claim? Seems to me, that's done by supposing that your claim is true in the same range of cases where your friend's claim is false, that your claim is false in the same range of cases where your friend's claim is true, and that your claim fails to have a truth value in the same intermediate range where your friend's claim has no truth value. Of course, we are then saying that in the intermediate range of cases, both you and your friend's claim match in that they both fail to have a truth value. But this seems to me the position that does the best job of making you contradict your friend. You seem to be denying the very same gappy thing that your friend is affirming. In any case where both of your claims have truth values, they have opposite truth values.

On the gap view, we get the same relation between our two speakers' claims in cases of contextually sensitive terms where there is a divergence in personally indicated content, where one speaker is asserting "P" and the other "Not-P". In the debate we have been imagining, our skeptic is denying the same gappy thing that her opponent is affirming.

What of the relation between the truth conditions of each speaker's utterance and the personally indicated content of that speaker? Here, the gap view does not deliver the closest match possible. Suppose you assert, "Frank knows that the Cubs won the 1908 World Series," where the epistemic standards you are personally indicating are different from – let's say lower than – those that your conversational partner is personally indicating, in a case where the two of you fail to adjust to one another's usage. We could suppose here that the truth conditions of your claim exactly match your own personally

indicated content – that your claim is true iff Frank meets your personally indicated standards and is otherwise false. Indeed, the multiple scoreboards view delivers this exact match. The gap view does not. On the gap view, in a certain range of cases – where Frank meets your, but not your partner’s, personally indicated standards – your claim is neither true nor false, though it would be true if the truth conditions of your claim simply followed your own personally indicated standards. But this difference is unavoidable on any view that, like the gap view, has you asserting exactly what your partner is denying. Where there is a divergence in personally indicated standards, no consistent view can make the truth conditions of both speakers’ claims match their own personally indicated standards while still delivering the result that one is denying precisely what the other is affirming.

However, while the gap view doesn’t do the best job possible of matching the truth conditions of each speaker’s claims with that speaker’s own personally indicated content, it does provide for what seems to me a nice relation between the two: In no possible case is your claim true where it would be false if it bore the content you personally indicate, and in no possible case is your claim false where it would be true if it bore the content you personally indicate. Alternative formulation: Wherever it is both the case that your claim has a truth value, and that it would have had a truth value if it bore the content you personally indicate, the truth value it has is the same as the truth value it would have had if it bore the content you personally indicate. Though, on the gap view, your claim can sometimes fail to have a truth value when it would have had a truth value if its truth conditions just were those that you personally indicate, it will never happen that your claim takes the opposite truth value from what it would have born if it had as its truth conditions what you personally indicate.

We might go so far as to call this relation that, on the gap view, the truth-conditional content of your claim bears to your personally indicated content that of “weak equivalence.” It’s the same relation that I think “P” bears to “It is true that P”. Where “P” is true, so is “It is true that P”; where “P” is false, so is “It is true that P”; but where “P” is neither true nor false, then it seems to me, as it seems to many others, “It is true that P” is false:

<u>P</u>	<u>It is true that P</u>
T	T
F	F
N	F

In no case is one of “P” and “It is true that P” true while the other is false, or, alternatively, wherever they both have a truth value, it is the same truth value. On the gap view, the truth-conditional content of your claims bears that same relation to your own personally indicated content.

The gap view achieves this, to my thinking, very impressive result of going a long way toward respecting both the sense of contraction and the match between the truth conditions of each speaker’s claim and that speaker’s personally indicated content only by assigning a gap in the truth conditions of the relevant claims, and in cases of great divergence in personally indicated content, this gap in the truth conditions – the range of cases where the claims are neither true nor false – gets very large. So the view is likely attractive only to those who are not great fans of bivalence. But I, for one, don’t mind gaps. And in very defective conversations, where there is great divergence in personally indicated content, I don’t find it implausible to suppose that those gaps get very large. In fact, that seems me the right verdict to reach about such conversations. After all, I found the exploding scoreboard view, on which the gap becomes all-engulfing, plausible as applied to cases of great divergence.

So, somewhat tentatively, I accept the gap view. On it, as on the exploding scoreboard view, neither the skeptic nor her opponent is speaking the truth in the situations we’ve been considering in this paper.¹⁰ I should report that when I asked Stewart Cohen about the situations under consideration, he too was inclined to say that neither the skeptic nor her opponent was speaking the truth, so he is probably inclined toward something like the gap view or the exploding scoreboard view.

Again, some may be worried about the extent to which the skeptic “wins” on these views, since, though she is not construed as speaking the truth when she claims her opponent does not “know,” she does succeed at creating a context in which her opponents cannot truthfully claim that she does “know”. But as I’ve indicated, I don’t much care to

what extent the skeptic “wins” in the sense in question. To my thinking, that’s neither here nor there.

Notes

¹ “Scorekeeping in a Language Game,” Journal of Philosophical Logic 8 (1979): 339-359; p. 355.

² It is tempting – indeed, almost irresistible – to describe the situation as one in which the skeptic is trying to raise the standards for knowledge, while her opponent is trying to keep the standards low. But since what I am here calling the standards for knowledge are the standards a subject must meet for her to be truthfully credited with “knowledge”, our debaters need not even be thinking of the standards, so construed, as being something that can be changed. They can be operating under the invariantist assumption that the truth conditions for knowledge sentences remain fixed. Thus, they shouldn’t be said to be trying to have an effect on the standards, so construed, which is why I instead described them as executing certain conversational maneuvers which have tendencies to put (or keep) in place certain standards. Typically, though, our debaters will be trying to change or affect other standards: the standards for when a subject will be accepted as “knowing” in their conversation, for instance. But they may conceive of themselves then as trying to get those latter standards in line with the “true” standards – those that articulate the conditions under which a subject can be truthfully said to “know” – as trying to get their opponent to admit the truth of the matter under discussion.

³ Here my use of the metaphor “scoreboard” differs from Lewis’s. For Lewis, there is a scoreboard “in the head” of each of the participants, and what the score is can be a function in part of what all these different scoreboards say the score is. As I use “scoreboard” here, it by definition gives the right score.

⁴ The need to stress the point of this paragraph was pointed out to me by Robert Stalnaker.

⁵ Due to his different use of “scoreboard,” described in the previous note, Lewis would not describe the position he tends toward as one in which there is a single “scoreboard”: He uses the term so that each speaker does have their own scoreboard in their head. However, Lewis does seem to assume in his writing that there is a single (though changeable) conversational score in a given conversation, rather than writing as if each there is a different score for each speaker, and thus seems to be working with the picture that in my use of the term has a single “scoreboard”.

⁶ “Scorekeeping in a Language Game,” Journal of Philosophical Logic 8 (1979): 339-359; p. 355.

⁷ One could quite plausibly understand Lewis differently here. For instance, if the force of Lewis's "must" is understood differently, this passage could be read as being compatible with the multiple scoreboards view that both the skeptic and her opponent are speaking the truth. Perhaps if Lewis's "commonsensical epistemologist" were to insist that he does "know," his claim would have the low, ordinary truth conditions, and would therefore be true, according to Lewis. The sense in which he "must" concede defeat, on this reading, is that if he is to speak properly, he must not make such an ordinary, low claim in response to the skeptic. Still, on this reading, a defiant response would be, though improper, true.

⁸ David Lewis, "Elusive Knowledge," Australasian Journal of Philosophy 74 (1996): 549-567.

⁹ There are ways to use "ignore" so that one can be said to be "ignoring" something that in some sense one is attending to: In a philosophical discussion, I am paying very careful attention to a certain objection that has just been raised, planning to craft a response later, but I am properly described as "ignoring" it as I proceed, in that I don't alter what I am now saying as a result of it. Similarly, if a skeptic brings to our attention a certain possibility of error, it seems to me that in a very good and relevant sense, we can choose to "ignore" it, even as we are thinking about it by, for instance, treating it as something not worth worrying about. But in "Elusive Knowledge," Lewis seems to me to use "ignore" in such a way that if you are at all attending to a possibility, you are thereby not ignoring it.

¹⁰ Throughout this paper, I have been considering only cases where just two speakers are involved in a conversation, and have been ignoring questions raised by cases where there are more than two speakers. Suppose, for instance, that there are 20 people, say, in a meeting of a philosophy seminar, discussing whether a certain character "knows" a certain fact. Suppose the seminar is not on the topic of skepticism, but rather concerns trying to formulate conditions for knowledge in light of Gettier examples. 19 of the 20 converge on personally adopting a certain set of very ordinary standards, while just one personally indicates other, very different, much higher standards. The one hold-out does not play a leading role in the discussion, making soft-spoken contributions ("Well, I think Henry doesn't know, because for all he knows, he's a brain in a vat"), which are scoffed at and otherwise ignored by others, only every so often. With so much agreement on a particular set of standards in the room, are we to suppose that this lone hold-out creates a situation in which the "knows" scoreboard explodes, or in which a huge gap is created in the truth conditions of everyone's claims? I'm not sure what to say, but it's tempting here to resort to a Lewisism (from "Scorekeeping," though Lewis uses it about a different kind of case) and say that if the character under discussion does meet the dominant standards in the room (the standards employed by the 19), then the claims of the 19 that the character does "know" are at least "true enough." Perhaps, though, the claims of the lone hold-out go gappy: When he claims that the character doesn't "know", his claim is true iff the character doesn't meet either his own personally indicated standards or the

dominant standards in the room; his claim is false iff the character meets both sets of standards; and if, as I'm supposing, the character meets the dominant standards in the conversation, but fails to meet our hold-out's personally indicated standards, the hold-out's claim is neither true nor false.