

Erik Stei

Epistemic Standards: High Hopes and Low Expectations

1. Epistemic Standards

Epistemic standards play an important role in many discussions about knowledge ascriptions: consider, for instance, the recent debate about the “factivity problem” for contextualism,¹ where Peter Baumann (2008) and Anthony Brueckner & Christopher Buford (2009) speak freely of contexts with “ordinary” standards and contexts with more demanding “sceptical” standards; Wolfgang Freitag (2011, 277), meanwhile, explicitly classifies epistemic standards, speaking of a “set of high-standard contexts” and a “set of low-standard contexts”, such that the union of both sets is the set of all contexts.² Of course, there may be a far greater variety of standards, but no matter how many different epistemic standards there are—two or indefinitely many—the question of what exactly it is that makes low standards low and high standards high needs to be addressed. We need a story about what it means to order contexts according to the strength of their respective epistemic standards. Given the importance of the notion for contextualism, subject-sensitive invariantism or epistemic relativism, it is surprising how few detailed and systematic accounts have been proposed given concerning how epistemic standards work. I discuss the most important ones in section 2.

In what follows, I presuppose a few things about epistemic standards that I take to be fairly uncontentious: First, epistemic standards can be raised. Maybe they can be lowered, too (Lewis 1996, 560), but this does not matter for my purposes here. Secondly, in the cases relevant to my discussion, the factors responsible for a rise in standards are “error-possibilities”, possibly but not necessarily accompanied by an increase in the practical importance of the proposition p of which the speaker is attributed or denied knowledge. Something along those lines is endorsed by all contextualists. Take, for instance, Keith DeRose: “The mentioning of alternatives like painted mules, or barn facades, or changes in banking hours [...] can be seen as raising the strength and changing the content

¹ Variants of the factivity problem are developed by Brendel (2005), Williamson (2001), and Wright (2005).

² The latter is a simplifying assumption. As far as I can tell, Freitag is happy to allow more than two standards.

of ‘know’” (DeRose 1992, 992). Similarly, Michael Williams claims that “raising and lowering of standards consists in the expansion and contraction of the range of error-possibilities in play” (Williams 2001, 2). Related characterisations can be found in Cohen 1999 or Lewis 1996. Thirdly, in stereotypical contextualist cases, contexts in which a given knowledge ascription comes out true have lower standards than contexts in which the same ascription comes out false.³ Any contextualist should happily subscribe to these assumptions.

Little work has been done on the details of epistemic standards. Exceptions are the views discussed in section 2, some critical comments on those views (e. g. Hawthorne 2004, Stanley 2005) and a paper by Jonathan Schaffer (2005) in which he argues that a “point-like” conception of epistemic standards in terms of relevant alternatives is plausible, whereas conceptions that rely on “thresholds” or “standards” are problematic.⁴ I agree with a lot of what he says, but I modify some of his results and generalize others. The general claims I defend in this paper are as follows: First, where the above assumptions are combined with a view that links epistemic standards to a measure on alternatives (including a contextually determined threshold) that allows the relevant set of contexts to be totally ordered, the resulting view leads to highly implausible results. Secondly, a theory that avoids these kinds of measures as well as a total ordering of contexts according to their epistemic standards has no resources to make sense of low or high standards. The conclusion I draw is that, as things stand, talk of low or high standards is ill-founded.

2. Three Models of Epistemic Standards

I discuss three ways of modeling epistemic standards that have been proposed by the most prominent defenders of contextualism (cf. Cohen 1988, DeRose 1995, Lewis 1979, 1996). According to those views, standards are associated either with modality, or with probabilities, or with quantifier domains. I briefly sketch each view in turn.

³ Note that this is not the same as presupposing that contexts have low and high epistemic standards *simpliciter*, i. e. standards that are just low or just high, as opposed to low or high compared to the standards of some other context.

⁴ Schaffer’s terminology differs from mine. My use of “standards” includes all of the conceptions Schaffer discusses. He uses “standards” to refer to what I call the “truth tracking” model or the “spheres” model, which is one way to model epistemic “standards” (in the sense in which I use the word).

Keith DeRose (1995, 33–38) links epistemic standards to modal considerations underlying his externalist view on the strength of the “epistemic position” of a subject. A subject S is in a strong epistemic position with respect to her belief p if her belief is not only true at the actual world, “but also at the worlds sufficiently close to the actual world” (DeRose 1995, 34), where closeness is measured in terms of similarity to the actual world.⁵ Now, generally, S knows p just in case her epistemic position is strong enough. Just how strong it needs to be in order to be strong enough for knowledge is determined by the epistemic standards of the context of ascription. Accordingly, DeRose invites us to picture his view of epistemic standards “as a contextually determined sphere of possible worlds, centered on the actual world within which a subject’s belief as to whether p is true must match the fact of the matter in order for the subject to count as knowing” (DeRose 1995, 36). A belief that is true in all worlds of a relatively small sphere may turn out to be a false belief in some worlds of a larger sphere. Thus the sphere symbolizes the epistemically relevant worlds and its extent is determined by the context in which the knowledge ascription is uttered. The mechanism of raising standards is given by the Rule of Sensitivity (DeRose 1995, 37): “When it’s asserted that S knows (or doesn’t know) that P , then, if necessary, enlarge the sphere of epistemically relevant worlds so that it at least includes the closest worlds in which P is false.” Accordingly, what counts as epistemically relevant depends on the contextually determined sphere. With the standards for knowledge linked to the extent of the sphere via epistemic relevance of alternatives, we get the following connection: The higher the epistemic standards, the larger the sphere of epistemically relevant worlds. Interestingly, as a belief that is true in all worlds of some sphere cannot be a belief that is false in some worlds of a smaller sphere, another result is that if the sphere associated with context c_1 is bigger than or as big as the sphere associated with c_2 , and if S “knows” p in c_1 , then S “knows” p in c_2 .⁶

Stewart Cohen (1988) agrees that context determines which alternatives are epistemically relevant, but he links relevance to an internalist probabilistic picture. Some not- p alternative h “is relevant, if the probability of h conditional on reason r and certain features of the circumstances is sufficiently high (where the

5 It is not plain overall similarity that is at issue here, but similarity with respect to the methods of belief formation. This takes care of scenarios like Nozick’s *grandmother case* (Nozick 1981, 179; DeRose 1995, 20).

6 As contextualism is in part a metalinguistic position about “knows”, I use quotation marks to indicate the dependence of the semantic value of “knows” on context. More precisely the phrase “ S ‘knows’ p in c ” is to be understood as “the utterance of the sentence ‘ S knows p ’ if made in c is true”. Thanks to an anonymous reviewer for pressing that point.

level of probability that is sufficient is *determined by context*)” (Cohen 1988, 103). In case the subject has sufficient reason to deny the alternatives with a probability above this contextually determined level, there are no relevant alternatives precluding her from knowing p . Accordingly, alternatives with a probability below a certain threshold are irrelevant for S 's knowledge of p in that context, even if they are compatible with the subject's reasons r . The connection we get on this picture is as follows: The higher the epistemic standards, the lower the probability of not- p alternatives (conditional on r) needs to be in order for those alternatives to be relevant. Like in DeRose's case, we also get a further result. As no not- p possibility that is below the threshold of some context can be relevant in less demanding contexts with higher thresholds, it follows that if the threshold associated with c_1 is lower than or equal to the threshold associated with c_2 , and if S “knows” p in c_1 , then S “knows” p in c_2 .

Finally, David Lewis prominently discusses epistemic standards on two occasions: the first when applying his conception of conversational score to relative modality (Lewis 1979), the second when developing his theory of knowledge ascriptions (Lewis 1996). In the first paper, Lewis classifies the dynamics of knowledge ascriptions as an instance of the more general connection between conversational score and modal expressions (see Lewis 1979, 354–355). He speaks of a “boundary between the relevant possibilities and the ignored ones”, which enters into the truth conditions of sentences containing expressions like “must”, “can”, and also “knows”. This boundary can be formally modelled by an accessibility relation between possible worlds, which may change in the course of conversation—that is, in the terminology adopted here, from context to context. Depending on the constraints on the accessibility relation, there are a number of ways to understand the workings of that boundary: for instance, if one relies on similarity between possible worlds (with respect to actual methods of belief formation), the view resembles DeRose's picture. Lewis's talk of attending to and (proper) ignoring of not- p possibilities, however, points to a view developed in more detail in his *Elusive Knowledge* (1996). In the definition of knowledge developed there, context-dependent standards enter in terms of properly ignored not- p possibilities. As every not- p possibility not properly ignored must be eliminated by the subject's evidence and what is properly ignored depends on context, the account can be linked to the context-dependency of quantifier domain restrictions (see also Ichikawa 2011). Here is the definition: “ S knows that p iff S 's evidence eliminates every possibility in which not- p [...] except for those possibilities that we are properly ignoring” (Lewis 1996, 554). So every not- p possibility that is not properly ignored is epistemically relevant.

Proper ignoring is tied to several rules which, among other things, are meant to take care of the factivity of knowledge (Lewis 1996, 554) and various epistemo-

logical problems like the lottery paradox and Gettier cases (Lewis 1996, 556). Most important for the purpose of this paper, however, are the Rule of Belief and the Rule of Attention as they reflect the contextual influence on epistemic relevance and, therefore, on epistemic standards. According to the Rule of Belief, no possibility is properly ignored “if the subject gives it, or ought to give it, a degree of belief that is sufficiently high” (Lewis 1996, 555), where what is sufficiently high may depend on how much is at stake. Lewis’s definition, as well as his examples (1996, 556), suggest that it is the stakes of the epistemic subject as well as objective stakes that matter. The rule may become sensitive to the context of ascription in case the ascribers know that the subject ought to give a higher degree of belief to a given not- p possibility than she actually does. The Rule of Attention is explicitly focused on the context of ascription (see Lewis 1996, 561). A not- p possibility attended to in the context of ascription may not be properly ignored. By attending to a possibility, ascribers make it relevant for a knowledge ascription in that context. If the subject’s evidence does not eliminate that possibility, the subject cannot be said to “know” that p in that context. The rule does not require a certain threshold, probability or degree. Attendance is an all-or-nothing affair.

It seems the Rule of Attention is more powerful with respect to ascriber sensitivity than the Rule of Belief. The cases in which ascribers know that the epistemic subject should give a higher degree of belief to a given not- p possibility are cases in which the ascribers attend to that possibility. This is sufficient to make that possibility epistemically relevant. Cases in which the ascribers or the subject fail to give a sufficiently high degree of belief to a possibility are cases in which S does not know that p for reasons independent of the context of ascription. So, in order to capture the contextualist cases at issue here, we need to focus on the Rule of Attention. As attendance is an all-or-nothing affair, there is no measure of relevance connected to alternatives on Lewis’s view. Thus, the only correlation we get on his picture is this: Given two contexts c_1 and c_2 , the epistemic standards of c_2 are higher than the epistemic standards of c_1 just in case the set of relevant alternatives of c_1 is a proper subset of the relevant alternatives of c_2 . We are not able to say that, generally, some not- p possibility q_1 is to be ranked higher, epistemically speaking, than some other not- p possibility q_2 . Given a context in which q_1 is salient but not q_2 , and another context in which q_2 is salient but not q_1 , we simply cannot tell whether the standards in the first context are higher than or equal to the other. On Lewis’s account, there is nothing that would allow us to speak of higher or lower standards in case the sets of possibilities relevant in c_1 and c_2 are disjoint or even in case they intersect without one being a subset of the other. But we do get the result that if the relevant alternatives of c_2 are a subset (proper or improper) of the relevant alternatives of c_1 , and if S “knows” p

in c_1 , then S “knows” p in c_2 . If S ’s evidence is incompatible with every element of a given set of alternatives, it is also incompatible with every element of a subset of that set.

It is worth noting that although some of the views discussed here allow for a general notion of epistemic strength (in terms of modal closeness or probabilities), none of the views offers any kind of general (i.e. context-independent) threshold, such that all contexts with epistemic standards above that threshold count as high standards contexts and all others as low standards contexts. So, except for the limiting cases in which either all alternatives are relevant or else no alternatives at all are relevant, it seems to make little sense to speak of absolutely high or absolutely low standards. Some views, however, do allow the notion of *relatively* high or *relatively* low standards, thereby enabling a comparative reading of epistemic standards: DeRose’s and Cohen’s views generally admit comparisons of contexts with respect to their epistemic standards (either in terms of modal distance or in terms of probabilities), while Lewis’s analysis only allows a rather restricted form of comparison.

3. Orderings of Epistemic Standards and Some Problematic Results

We are now in a position to take a closer look at the way contexts may be ordered according to their epistemic standards. Assume a set M_c containing contexts of utterance. As we will see, there may or may not be restrictions as to which contexts may be elements of M_c . The comparative reading of epistemic standards in the spirit of Cohen’s and DeRose’s views requires a binary relation “ \leq ” on the elements of M_c (such that $\leq \subseteq M_c \times M_c$) that allows at least a partial ordering. If “ $c_1 \leq c_2$ ” is to represent the intuitive notion that the standards of c_1 are lower than or equal to the standards of c_2 , “ \leq ” needs to be reflexive, transitive and antisymmetric.⁷

It was shown that DeRose’s sphere model links epistemic standards to the extent of a sphere of possible worlds. As the worlds are centred on the actual world and ordered by a similarity measure, we can interpret “ $c_1 \leq_s c_2$ ” as “the sphere of possible worlds of c_1 is smaller than or equal to the one of c_2 ”.⁸ As

⁷ Thus, for any $a, b, c \in M_c$: $a \leq a$; if $a \leq b$ and $b \leq c$, then $a \leq c$; and if $a \leq b$ and $b \leq a$, then $a = b$.

⁸ The additional restriction on the method of belief formation—only worlds in which S forms her belief that p in the same way she does in actuality are to be considered—leads to a constraint on M_c . Worlds in which S forms her belief in different ways are not epistemically relevant, as they

we have seen in the previous section, contexts are ordered according to how far the sphere of epistemically relevant worlds reaches into modal space. According to this reading, M_c appears to be totally ordered, as it meets the additional constraint that for any $a, b \in M_c$: $a \leq b$ or $b \leq a$.⁹ The reason for this is that (with the similarity ordering of possible worlds remaining fixed) the only thing that varies with context is the extent of the sphere. Given a contextually determined sphere s_1 encompassing all worlds up to distance d_1 , any other contextually determined sphere s_2 encompassing all worlds up to distance d_2 will either be as large as s_1 or smaller than s_1 or larger than s_1 . For every pair of members of M_c , one is smaller than or equal to the other. This allows for a contextually determined value, such that if S “knows” p in c_1 , then S “knows” p in all contexts c_n for which $c_n \leq_S c_1$.

This leads us to the first counterintuitive result: the problem with this picture is that shifts in epistemic standards “globally infect” other propositions believed by the epistemic subject (see Schaffer 2005, 124). Given DeRose’s conception of the closeness relation between possible worlds, it is an immediate consequence that even propositions unrelated to p and known in a low-stakes context may not be known in a high-stakes context. The reason is that modal closeness is measured by similarity to the actual world and the actual method of how S ’s belief that p is formed. If standards are raised, the sphere of epistemically relevant worlds includes *all* worlds up to the contextually determined similarity value. In some of those worlds, propositions other than p which are truly believed by the subject in actuality will turn out to be falsely believed by the subject. Thus, the epistemic subject does not know those propositions in the high-stakes context, as her epistemic position is not strong enough: her belief is not true in all worlds sufficiently close to the actual world. This may be a suitable position with respect to radical scepticism—the problem DeRose is dealing with when developing his view. Should the possibility of S being a brain in a vat be uttered, S loses “knowledge” of a considerable number of propositions, at least according to the contextualist view. It is far less convincing with respect to more moderate kinds of epistemic doubt. Consider, for instance, Cohen’s airport case. Passenger

cannot be pictured as being inside or outside of any sphere centred on actuality including S ’s method of belief formation.

⁹ Schaffer (2005, 124) proposes a different view according to which the conversational context not only determines the size of the sphere but also the way in which the possible worlds are ordered (Schaffer’s metric m). This results in a very powerful conception of context. My impression is that the above reconstruction, where the objective features of the world of the subject’s context determine the ordering (but not the sphere), is closer to DeRose’s view. Nothing of importance for the present paper depends on this, however, as the cases considered here are such that in both contexts the way in which possible worlds are ordered remains fixed.

Smith is asked whether the plane stops in Chicago. After looking at his flight itinerary, he responds “Yes, I know it stops in Chicago”. Assume that in this context the sphere of possible worlds is of size n . After the possibility of a misprint in Smith’s itinerary is brought up, context is changed by a rise in epistemic standards and the sphere becomes larger, now measuring $n + 1$. It is not only Smith’s belief that the plane stops in Chicago that is false in some worlds in the modal space between n and $n + 1$, but many other beliefs as well. So Smith not only loses “knowledge” of the proposition that the plane stops in Chicago; “knowledge” of propositions completely unrelated to p may be lost too, namely all those propositions that Smith believes truly in n , but that he falsely believes in some worlds between n and $n + 1$.

I take this to be a highly implausible result. Suppose that the gate for the flight in question is not mentioned in the itinerary, but Smith reads it on the destination board. Given that worlds in which the itinerary contains a misprint are at least as far off in terms of similarity to actuality as worlds in which the gate of the flight is changed (a plausible assumption, I think), why should mentioning the possibility of a misprint exclude Smith’s “knowledge” that the flight leaves at gate 46? If we suppose that it is both true that the plane stops in Chicago and that the departure is at gate 46, and that Smith is in a reasonably good epistemic position with respect to those propositions, such that he knows both of them in the lower-standards context, it seems false to say that Smith loses “knowledge” that the departure is at gate 46 just because an error-possibility as to whether the itinerary might have been changed is raised. However, this is exactly what the sphere model predicts.

A similar problem can be constructed for Cohen’s conception of epistemic standards. We saw in the last section that, on Cohen’s view, the higher the epistemic standards, the lower the probability of not- p alternatives needs to be in order for those alternatives to be relevant. Thus, we can read the relation “ $c_1 \leq_T c_2$ ” as “the (probability) threshold of c_1 is lower than the threshold of c_2 ”, resulting in a total ordering on M_c with decreasing strictness of epistemic standards. On this view, M_c is totally ordered, just like in DeRose’s sphere model. For every pair of members of M_c , the threshold of one is smaller than or equal to the other’s. This allows for a contextually determined value such that if S “knows” p in c_1 , then S “knows” p in all contexts c_n for which $c_1 \leq_T c_n$.

One way of interpreting Cohen’s suggestion is that context influences the threshold for how probable (conditional on S ’s reason r) an alternative to any of S ’s belief must be to count as a relevant alternative. On this reading, discussed by Schaffer (2005, 118–121), the probability model is just as globally infectious as the sphere model. Mentioning error-possibilities with respect to the belief that p or raising the practical importance of p would globally lower the proba-

bility for alternatives to S 's beliefs, stripping her of "knowledge" of all propositions with alternatives relevant according to the new probability threshold. Thus, counterexamples like the gate/misprint case discussed above would arise just as easily.

But there is another way to reconstruct the position. Unlike the sphere model, the probability model explicitly refers to not- p alternatives. So a sensible option may be to suppose that only the threshold value for the probability of not- p alternatives (conditional on r) shifts, while the threshold for possibilities unrelated to p may remain untouched. This may affect the actual ordering, but not the structure of the relation " \leq " on M_c , however. According to this picture " $c_1 \leq_T c_2$ " would have to be read as "the threshold for not- p possibilities in c_1 is lower than the threshold for not- p possibilities in c_2 ", still resulting in a total ordering of the contexts of M_c according to the probability of the alternatives relevant in a context. But it would render the model less infectious than the sphere model. In case Smith cannot deny the alternative that the itinerary contains a misprint, this may prevent him from "knowing" that the plane stops in Chicago, but it need not prevent him from "knowing" that the departure gate is 46 even if the probability of the latter possibility is equal to or higher than the first. Possibilities in which the gate is changed are epistemically independent of possibilities in which the flight does not stop in Chicago. A change of the probability threshold for not- p possibilities need not affect all of S 's other beliefs.

This may be an improvement, but the position is problematic nonetheless. The focus on not- p possibilities is still too broad, as it leads to other counterintuitive cases. Imagine two possibilities: Smith's itinerary contains a misprint (q_1), and the plane in question suffers an engine defect when steering out to the manoeuvring area, resulting in the cancellation of the flight (q_2). Let us suppose for the sake of argument that both possibilities have the same probability $0.n$ (conditional on r). Let us assume, further, that both q_1 and q_2 are not- p possibilities. Now imagine, as before, that the plane really stops in Chicago, but that the possibility of a misprint in Smith's itinerary is raised. Thus, the probability threshold for not- p alternatives is lowered accordingly to value $0.n$. Now, Smith decides to check back at the counter of the flight company and has his itinerary confirmed. He cannot, however, deny the possibility that the flight is cancelled because of an engine defect, which also takes value $0.n$. Now, given that the possibility of a defective engine is as probable as the possibility of a misprint, he cannot be said to "know" that the plane stops in Chicago (in that context), because the proba-

bility threshold on not- p possibilities is too low (in that context). Again, this is a highly counterintuitive result.¹⁰

One might object that an error-possibility, in that case q_2 , needs to be salient in order to be relevant in the first place, but we have seen that, given \leq_T , and irrespective of the specific not- p possibilities of a context, if S “knows” p in c_1 , then S “knows” p in all contexts c_n , such that $c_1 \leq_T c_n$. Besides, the insistence on salience immediately leads to further problems. One I will not address here is how to avoid knowledge ascriptions coming out as true far too easily in contexts of ignorant subjects and/or ascribers. More important for my purposes is this: the only sensible use the probability of an alternative (conditional on r) as an additional criterion to salience could have is that some salient alternatives may be ignored, namely those that are irrelevant in that context according to the probability threshold. The problem is that on practically all contextualist accounts (and certainly on the ones discussed so far), salience is usually sufficient to raise epistemic standards. So what would be needed is an account of how salience sometimes does and sometimes does not raise the standards. It is not at all clear what this account would look like.¹¹ Besides, why should an ordering be introduced in the first place, if salience does all the important work?

In the last section, it was shown that the only cases in which we can speak of higher or lower standards on Lewis’s picture is if the set of relevant not- p possibilities of one context is a proper subset of the set of relevant not- p possibilities of another context. It is not possible, however, to compare the epistemic standards of contexts in which this subset requirement is not fulfilled. So there are pairs of contexts for which we cannot say which one has the higher standards or, more precisely, it is not the case that for any element a, b of M_c , $a \leq b$ or $b \leq a$. So, although it also holds on Lewis’s view that, for all contexts, if $c_n \leq c_1$, and if S “knows” p in c_1 , then S “knows” p in c_n , it is impossible for c_n to contain relevant alternatives that are not also in c_1 , due to the restrictions on the relation “ \leq ”. This seems to be exactly the feature that avoids the counterexamples

10 Note that Cohen’s analysis of the lottery case (Cohen 1988: 106–108) does not help to avoid the problem: as the chance of error is the same for both q_1 and q_2 , it does not matter whether it is salient or not. Further, it is one and the same context we are considering, so Smith’s reasons for his belief that the plane stops in Chicago remain unchanged. Besides, it is not a radically sceptical result: possibilities that are less probable (conditional on r) than q_1 and q_2 remain epistemically irrelevant.

11 Cohen notes (1999, 85 n. 27) that salience does not necessarily raise the standards, but he does not offer a systematic analysis of the phenomenon. To be sure, there are accounts that deal with this problem (see, e.g., Blome Tillmann 2009), but those are typically ones that do not picture contexts as totally ordered according to their epistemic standards.

against DeRose's and Cohen's conceptions of epistemic standards. As there is no measure on not- p possibilities, the only context-dependent feature affecting their epistemic relevance is *being attended to*. So there is no context-dependent way a not- p possibility not attended to could become relevant.¹² Just because a misprint in Smith's itinerary is relevant (because of its salience), there is nothing in Lewis's account that automatically makes other possibilities relevant. The above counterexamples can be avoided.

But this comes at a cost. In the absence of a measure that would allow us to order contexts according to their epistemic standards, it seems that we cannot speak of high and low standards or of "everyday" and "sceptical" standards in an absolute sense. Even talk about standards of one context being higher or lower than the ones of another context is quite limited with a Lewis-style partial ordering. Even if we know, for instance, $c_1 \leq c_2$ and $c_3 \leq c_4$, it may be impossible to say anything about the relation of the epistemic standards of c_1 and c_4 . Take, for instance, $M_x = \{c_1, c_2, c_3, c_4\}$ and the partial ordering $\leq_x = \{<c_1, c_1>, <c_2, c_2>, <c_3, c_3>, <c_4, c_4>, <c_1, c_2>, <c_3, c_4>\}$. Given \leq_x , it is not only impossible to tell which contexts have low epistemic standards and which have high epistemic standards: there are also contexts for which we cannot even say whether their standards are higher or lower than those of other contexts. To make this point more vivid, imagine that \leq in the example above orders individuals according to their height. If for all we know c_1 is smaller than or as tall as c_2 and c_3 is smaller than or as tall as c_4 , then we are at a loss when asked whether e.g. c_1 is smaller than or as tall as c_4 . There is a consistent extension of \leq according to which c_1 is smaller than or as tall as c_2 but taller than both c_3 and c_4 , but there is also one according to which c_4 is taller than c_1 . The situation is similar with respect to Lewis's conception of epistemic standards. While avoiding the counterexamples, it severely limits talk of high or low epistemic standards.

4. Diagnosis and Conclusion

Abstracting a bit from the modified airport case, the counterexamples seem to build on a simple recipe: Take some not- p possibility q_1 that is incompatible with S 's epistemic situation (e.g. the possibility that Smith's itinerary contains a misprint, which turns out to be incompatible with Smith's evidence after he has checked back with the airline agent in the example above). Assume a meas-

¹² Of course, a possibility not attended to may be relevant due to other rules, but those are of no use when analysing the contextualist cases, so we need not worry about them here.

ure F assigning values to alternatives, such that q_1 takes value n . Take another not- p possibility q_2 with the same or a lower value as q_1 on F that is compatible with S 's epistemic situation (e.g. the possibility that the plane in question suffers an engine defect when steering out to the manoeuvring area, resulting in the cancellation of the flight). Describe a context c_1 where the epistemic standards are at level n at which neither q_1 nor q_2 are relevant and S "knows" p . Then describe a context c_2 resulting from c_1 where q_1 is brought to the ascriber's attention (e.g. the mentioning of the misprint possibility resulting in Smith's checking with the airline). This leads to a rise in epistemic standards to level $n + 1$, making q_2 relevant too (both q_1 and q_2 take the same value on F). Then, although S 's evidence is incompatible with the not- p possibilities salient in context c_1 (including q_1), she still does not know p , because her evidence is not inconsistent with every alternative relevant in c_1 according to F . After all, q_2 is consistent with her evidence by hypothesis.

Now, what is the fundamental problem this recipe takes advantage of? We have seen that, in the case of DeRose and Cohen, a measure of epistemic strength formulated in terms of the epistemic relevance of alternatives resulting in a total ordering of contexts according to their epistemic standards leads to counterexamples. Is it enough to reject one of those two characteristics, that is, either the measure on epistemic strength or the total ordering, in order to avoid the problematic cases? As far as I can see, the dependencies are as follows: First, if contexts are totally ordered, even absent any measure for epistemic standards it will be possible to construct counterexamples—at least if the ordering respects the intuitively plausible constraint that if $c_n \leq c_1$, and if S "knows" p in c_1 , then S "knows" p in c_n .¹³ The recipe will have to be rephrased such that it does not refer to values on F , but this is only a change in detail, not in spirit. Describe contexts c_1 and c_2 in the way outlined above with the modification that the epistemic standards of c_2 are lower than or equal to c_1 given the relevant relation " \leq " on the set of contexts (instead of the measure F).

Further, we have seen that Lewis's account is not affected by the counterexamples as it dispenses with a measure on epistemic strength and imposes only a partial order (in combination with the constraint that if $c_n \leq c_1$, and if S "knows" p in c_1 , then S "knows" p in c_n). But this comes at the cost of seriously restricting the discourse about epistemic standards. Is it possible to improve on this? It seems that Lewis's account could only allow a more extensive discourse about standards if the subset requirement were dropped and an independent measure

¹³ This constraint immediately followed from Cohen's and DeRose's accounts, but it might have to be stipulated as an additional constraint on alternative conceptions.

for the strictness of epistemic standards added. What would be needed, in other words, is some F_x that assigns values to the alternatives relevant in a context, such that, irrespective of the details of those values, the resulting ordering on the class of contexts M_x will not be total. In light of the results above it seems unlikely that this will be helpful. Assume, for the sake of an example, that M_x contains four contexts $c_1, c_2, c_3,$ and c_4 with F assigning the values $\{1, 2, 4, a\}$ to contexts in view of the epistemic strength needed for the subject's evidence to be strong enough for "knowledge" (with $1 < 2 < 4$ and " a " symbolizing something like "unsettled" standards¹⁴): $c_1 = 1, c_2 = 2, c_3 = a, c_4 = 4$. The resulting ordering relation would be as follows: $\leq_x = \{<c_1, c_1>, <c_2, c_2>, <c_3, c_3>, <c_4, c_4>, <c_1, c_2>, <c_1, c_4>, <c_2, c_4>\}$. This would allow us to say, for instance, that standards in c_4 are higher than in both c_1 and c_2 , but the ordering would not be total because no context except c_3 itself stands in relation " \leq_x " to c_3 . But it would also allow for counterexamples following the recipe above. Let q_1 and q_2 be not- p -possibilities with value 4. Assume that q_1 is incompatible with S 's epistemic situation while q_2 is compatible with it. Imagine the ascribers are in c_1 so neither q_1 nor q_2 are relevant. Assume further that S "knows" p in c_1 . Now assume that q_1 is uttered, causing the epistemic standards to rise to value 4. Will q_2 (and any other not- p -possibility with value $n \leq 4$) be relevant in the resulting context? If one wants to keep the constraint that if $c_n \leq c_1$, and if S "knows" p in c_1 , then S "knows" p in c_n , one is compelled to say "yes", but then S does not "know" p in c_4 , as there is an alternative, namely q_2 , compatible with her evidence. If one refuses to make every alternative with a value $n \leq 4$ relevant, one will have to give up on the constraint: S may be said to "know" in c_4 but not in a different context c_{4^*} which differs from c_4 only insofar as q_2 (and all other not- p -possibilities with value $n \leq 4$) is relevant as well. As both contexts will be "value 4"-contexts, the fact that S "knows" p in c_4 but S does not "know" p in c_{4^*} violates the constraint.

If these considerations are correct, then in order to avoid the counterexamples it is necessary to avoid both a total ordering of contexts according to their epistemic standards and an independent measure F that assigns values to alternatives. A further option may be to give up the constraint that if $c_n \leq c_1$, and if S "knows" p in c_1 , then S "knows" p in c_n , but this means giving up a lot of the initial attraction of contextualism. If knowing in high standards, like for instance sceptical standards, has no consequences whatsoever for knowledge in lower standards, like everyday standards, then the notion of standards loses its pull. The result is that either the intuitively plausible talk of (relatively) low- and (rel-

14 For the notion of "unsettled standards" see, e.g., Montminy and Skolits 2014.

actively) high-standard contexts must be abandoned, or its defenders must find another way to block the counterexamples.¹⁵

References

- Blome-Tillmann, M. (2009) Knowledge and Presuppositions. In: *Mind* 118, 241–294.
- Brendel, E. (2005). Why Contextualists Cannot Know They Are Right: Self-Refuting Implications of Contextualism. In: *Acta Analytica* 20(2), 38–55.
- Cohen, S. (1988). How to be a Fallibilist. In: *Philosophical Perspectives* 2, 91–123.
- Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reason. In: *Noûs* 33, 57–89.
- DeRose, K. (1992). Contextualism and Knowledge Attributions. In: *Philosophy and Phenomenological Research* 52(4), 913–929.
- DeRose, K. (1995). Solving the Skeptical Problem. In: *The Philosophical Review* 104(1), 1–52.
- DeRose, K. (2004). Single Scoreboard Semantics. In: *Philosophical Studies* 119, 1–21.
- Freitag, W. (2011). Epistemic Contextualism and the Knowability Problem. In: *Acta Analytica* 26, 273–284.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Ichikawa, J. (2011). Quantifiers and Epistemic Contextualism, In: *Philosophical Studies* 155, 383–398.
- Lewis, D. (1979). Scorekeeping in a Language Game. In: *Journal of Philosophical Logic* 8(3), 339–359.
- Lewis, D. (1996). Elusive Knowledge. In: *Australasian Journal of Philosophy* 74(4), 549–567.
- Montminy, M. and Skolits, W. (2014). Defending the Coherence of Contextualism. In: *Episteme* 11, 319–333.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge.
- Schaffer, J. (2005). What Shifts? Thresholds, Standards, or Alternatives? In: Preyer, P. and Peter, G. (eds.), *Contextualism in Philosophy. Knowledge, Meaning, and Truth*. Oxford, 115–130.
- Stanley, J. (2005). *Knowledge and Practical Interest*. Oxford: Oxford University Press.
- Williams, M. (2001). Contextualism, Externalism and Epistemic Standards, In: *Philosophical Studies* 103: 1–23.
- Williamson, T. (2001). Comments on Michael Williams’ “Contextualism, Externalism and Epistemic Standards”. In: *Philosophical Studies* 103, 25–33.
- Wright, C. (2005). Contextualism and Scepticism: Even-Handedness, Factivity and Surreptitiously Raising Standards. In: *Philosophical Quarterly* 55, 236–262.

¹⁵ I would like to thank the audience of the workshop “Neue Perspektiven in der Erkenntnistheorie II” at TU Dresden for discussion and comments on the material presented in this paper. Thank you also to an anonymous referee and especially to Martin Grajner and Pedro Schmechtig for their patience and for their helpful reactions to an earlier draft of this paper.