

Varieties of Contextualism: Standards and Descriptions¹

Peter Baumann

University of Aberdeen

Summary

Most contextualists agree that contexts differ with respect to relevant epistemic standards. In this paper, I discuss the idea that the difference between more modest and stricter standards should be explained in terms of the closeness or remoteness of relevant possible worlds. I argue that there are serious problems with this version of contextualism. In the second part of the paper, I argue for another form of contextualism that has little to do with standards and a lot with the well-known problem of the reference class. This paper also illustrates the fact that contextualism comes in many varieties.

Only a couple of years ago, one could still easily make sense of general questions like “What do you think about epistemological contextualism?” In the meantime, so many different positions have been developed under the heading of “contextualism” that one is tempted to reply “It depends on what you mean by ‘contextualism’”. The variety has become so great that what is a serious objection to one form of contextualism might be even welcome support for another form of contextualism. Therefore, one should look at them one by one. One distinction is that between attributor and subject contextualism (cf. DeRose 1992, 918ff.; DeRose 1999, 190f.; Cohen 1987; Cohen 1988). In the meantime, it has become more popular to refer to the latter under the heading of “subject sensitive invariantism” (cf., e.g., Hawthorne 2004, ch.4) and to reserve the term “contextualism” for the former. I will go with this and take contextualism to be, broadly speaking, the following thesis: The truth-value of knowledge ascriptions of the

¹ Published in *Grazer Philosophische Studien* 69, 2005, 229-245.

form “S knows that p” (and of related forms, of course) may (but need not) change with the speaker’s context (or the thinker’s context, for that matter). That is, it may change from speaker to speaker or between different contexts one and the same speaker finds herself in. To choose what is perhaps the most overused example in this context: In one (an ordinary) context it might be true to say or think that Jack knows that he has hands, but in another (sceptical) context it might not be true.²

But what is a context and how should we individuate contexts? Most contextualists seem to agree that contexts differ with respect to the relevant epistemic standards. In an ordinary context, epistemic standards are modest enough to allow for Jack to know that he has hands; in a sceptical context, however, standards are so strict that he won’t even know that. So, contexts are determined by the epistemic standards which count as the relevant ones in a given situation. It has become so much standard among contextualists to make contexts depend on epistemic standards that one could also speak of “standard”-contextualism in both senses of the word. In what follows I will take a look at a very prominent way to explain the difference between more modest and stricter standards, and will also say why I’m not happy with it. After that, I will try to make a constructive alternative proposal and defend another form of contextualism – a version that has little to do with standards and deserves much more attention than it has received so far. But first to standards.

² I will take the stylistic freedom to switch from the metalinguistic level to the object-level whenever nothing hinges on it (cf. Lewis 1996, 567).

1. Standards

According to authors like DeRose 1995 and others (for a slightly different view cf. Cohen 1998, 289-291 or Cohen 1999, 57-60) S's knowledge that p requires that S can rule out or that S's evidence eliminates alternate possibilities incompatible with the truth of "p". To rule out a possibility, S needs to have evidence to the effect that the possibility in question is not actual. An uneliminated alternate possibility is one in which the subject would have exactly the same experience and evidence as they actually have (cf. Lewis 1996). For lack of a better term, I will say that S "eliminates" a possibility just in case S can either rule it out or S's evidence eliminates it. Now, which alternate possibilities does the subject need to eliminate? Well, according to the philosophers just referred to, that depends on the epistemic standards relevant in the context of the attributor (speaker or thinker). According to laxer standards, Jack has to eliminate fewer and less remote possibilities in order to make "Jack knows that there is a hot dog in front of him" true. For instance, he just has to eliminate the possibility that there is nothing but a fake hot dog, made out of wax, in front of him. According to stricter standards, Jack has to eliminate more and more remote possibilities in order to make "Jack knows there is a hot dog in front of him" true. For instance, he would have to be able to eliminate the possibility that he is a Cartesian dreamer dreaming of some hot dog in front of him when there is none. Given stricter standards, Jack has to eliminate all possibilities he has to eliminate under laxer standards plus some more remote possibilities.³ Given laxer standards, it is true to say that Jack knows that there

³ It seems to be a (usually implicit) assumption of many contextualists that if S can eliminate a more remote possibility, then they can *ipso facto* eliminate a closer possibility. It is not obvious at all whether this is the case. Is it not conceivable that I can eliminate the possibility that I am a robot from

is a hot dog in front of him, at least if we assume that Jack can tell wax dogs from real hot dogs. However, given stricter standards, this might not be true. One of the main points of contextualism is that there is no single “true” standard or context. There is rather a plurality of contexts and in some it is true whereas in others it is false to say that S knows that p. In this sense – but in this harmless sense only – contextualism is a form of relativism.

There are several general objections one could make against the version of contextualism just sketched (and I really only want to talk about the basic idea here and neglect finer differences between different versions of it). First, one might complain about a certain circularity in talk about “eliminating possibilities”. Suppose we claim that Jack knows there is a hot dog in front of him and that he can eliminate the possibility that it is only a wax imitation of a hot dog. If “eliminating the possibility that it is an imitation” means something like “knowing that it is no imitation”, then we are explaining “knowledge that there is a hot dog” in terms of “knowledge that it is not something else”. We are using the explanandum in the explanans and that is not good. What if “eliminating the possibility that q” does not mean or imply “knowing that not q”? Well, what could it mean then? Something like “having good reasons to believe that not q”? But even good reasons can be misleading. Jack has excellent reasons to believe there is no wax hot dog in front of him; however, unbeknownst to him, he is confronted with the latest high tech wax imitation of a hot dog and cannot tell it from a real hot dog. He therefore doesn’t know there is a hot dog in front of him. In other words, it seems we are facing a dilemma: Either we are reading “eliminating” in a sense that is too weak (e.g., as “having good reasons”); or we are reading it in a strong

Mars (remote possibility) but that I cannot, for some reason, eliminate the possibility that I was born on the 5th of May rather than, say, the 4th of May at 11.45 pm?

enough sense but then the situation is even worse because we smuggle the explanandum into the explanans. Hence, one might conclude, all this stuff about “eliminating possibilities” does not lead anywhere (on “ruling out” see also Dretske 1970 and 1981b).

The short reply to this objection is that we need not be interested at all in a reductive explanation of the concept of knowledge here (who has ever managed to reductively define concepts like that?). Apart from that, the objection only seems to work for the internalist interpretation of “eliminating a possibility” as “ruling it out”. The objection does not seem to work for, say, Lewis’ more externalist talk about “eliminating possibilities”. Hence, the objection above is, one might reply, without force.

I think one can reformulate the objection but there is a much more serious problem with the contextualist story about eliminating possibilities. It concerns the underlying assumption that one can rank possibilities or possible worlds, for that matter, in terms of their remoteness from the actual world (cf., e.g., DeRose 1995, 36-37 and, as a background, Nozick 1981, 172ff. on sensitivity; cf. also, in general: Williams 2004). Our contextualist wants to say that in a stricter context the subject has to eliminate all the possibilities she has to eliminate in laxer contexts plus more remote possibilities. It is false to say “Jack knows that he has hands” if he can only eliminate the possibility that he is a brain in a vat but not the possibility that he has stumps instead of hands.⁴ If there is something wrong with the idea of a ranking of possible worlds in terms of remoteness, then this version of contextualism collapses, as we will see. So, what could be wrong with the idea of a ranking?

4 This does not contradict the last footnote: The fact that knowledge requires that one can eliminate both the more remote and the closer possibility does not imply that one can indeed eliminate both.

Take the usual question whether Jack knows that he has hands. Let us assume that Jack has (two) hands. According to the kind of contextualism under consideration, it is true to say “Jack knows that he has hands” if the attributor finds herself in an ordinary context in which the standards only demand that the subject be able to eliminate close possibilities (like the one that he’s been in a car accident and lost both hands, and similar horror scenarios). However, it is false to say “Jack knows that he has hands” if the attributor finds herself in a sceptical context in which the standards also demand that the subject be able to eliminate remote possibilities like the one that he is just a handless brain in a vat who is merely imagining that he has got hands. The whole question I want to raise is this: Why should one believe that a brain in a vat-world is a remote possible world (cf. also Neta 2003, 16, fn. 51 who raises this question *en passant*)? Why believe a world in which Jack was involved in an accident leading to the loss of his hands is closer to the actual world (or to what we take to be the actual world)?

Usually, the closeness or remoteness of possible worlds is spelled out in terms of degrees of similarity between the worlds considered (cf. Lewis 1973, 48-52, 66-67; Lewis 1986, 20-27). Let us call a world in which Jack is involved in an accident an “accident-world” and a world in which he is just a brain in a vat a “vat-world”. Why should one think that an accident-world is more similar to actuality than a vat-world is? Compare an accident-world with very different laws of nature with a vat-world with our laws of nature. Which one is more similar to the actual world? Ask a physicist and he’ll probably tell you (if he is willing at all to engage in that kind of speculation) that the vat-world is more similar. What matters to him is not so much whether he or Jack or we all are envatted but what the correct physics for the world under consideration is.

Then ask an epistemologist and he'll probably give you the opposite answer. What matters to him is not so much the laws of nature but what his or Jack's or our epistemic situation is.⁵

So, ironically it all seems to depend on the context. Epistemologically, the accident-world is closer but ontologically the vat-world is closer to the actual world. A case like this suggests that there is no remoteness-ranking of worlds as such. Remoteness or closeness of possible worlds is relative to a context in which different standards of importance (physics, epistemology) are the ones that count. There is no context-independent hierarchy of contexts or standards in terms of strictness or laxness (cf. also Williams 1996 and Heller 1995, 505-507).⁶

The irony is that the contextualist story of people like DeRose and others is blocked by an unwelcome context-dependency.⁷ If there is a plurality of equally acceptable but mutually incompatible rankings of possible worlds in terms of remoteness and closeness, then we might, for instance, have to deal with two rankings like the

5 The physicist's perspective is a third person perspective whereas the epistemological perspective is a first person one: It is about "my" epistemic situation. I won't pursue things in terms of this aspect of the difference here. By the way: It seems obvious that there are many epistemologists who would and do in fact respond to the questions above in the way I indicated. Physicists seem less interested than some epistemologists in ranking possible worlds but as far as I can see they would also rather tend to respond in the way indicated. Just ask one! In the end, however, that does not matter too much for the argument. What matters is that nothing having to do with the characteristics of the different possible worlds gives us a reason to answer the above questions about closeness one way or another.

6 DeRose 1992, 922 (fn.18) points out (following Unger 1986) that we have different epistemic standards for different aspects of beliefs: for the degree of confidence, for the degree of non-accidentality of the belief, etc. In some contexts the standards for the necessary degree of confidence are very strict and the standards for non-accidentality are rather relaxed whereas in other contexts it is the other way around. Hence, we should not expect a unique hierarchy of contexts. Whatever one would have to say about this point, it is obvious that it is quite different from the point made in the text above which deals with just one aspect or dimension, namely the closeness or remoteness of possible worlds.

7 All this should not come as a surprise if one explains closeness in terms of similarity of worlds, like Lewis. The notion of similarity is notoriously vague, multi-dimensional and context dependent (cf. Fine 1975, 451-458, Jackson 1977, 4-8, Slote 1978, 20-25 and Heller 1999, 116; see also Goodman 1972 and Tversky 1977). Lewis himself accepts the context dependency of the notion of similarity (cf. Lewis 1973, 91-95) and thus of the closeness of possible worlds (cf. Lewis 1973, 50-52, 66-67). Jackson, Slote and others offer alternative accounts of counterfactuals but I doubt these accounts can make the above problem go away: It is a very general one.

following ones: one ranking according to which the possibility that I am a brain in a vat (or a robot from Mars) is less remote than the possibility that I have either 1 or 3 hands, and another, second ranking which gives the reverse order. Given such a plurality of rankings, the contextualist story against the sceptic seems to collapse: We cannot “secure” an ordinary context from sceptical threats anymore because one could always see the latter as very close to home (too close for a contextualist response to the sceptic). Apart from that, and more generally, the contextualist account of how the truth of knowledge claims varies with context shifts would lose its plausibility if spelled out in terms of rather arbitrary remoteness-rankings.

One way to respond to this would be to propose to take the second context-dependency on board and thus “deepen” and “radicalize” contextualism. What we then get is a “two-step”-contextualism. First, whether it is true to say that S knows that p depends on how far out the possibilities are that S has to eliminate, according to the attributor. Second, what is far out or really close depends on standards of relevance used by the attributor (not the subject). There are two dimensions of context-dependency and the latter one is more fundamental than the former one (the one DeRose and others have in mind). In the end, it still all depends on the attributor’s standards of relevance.

I am sceptical about this kind of reply. It rescues the DeRose-style of contextualism by transforming it into something different. This move still seems to make things a bit arbitrary: It all depends on what an attributor happens to think is relevant. You might think this is relevant and I might think that is relevant, and that seems to be the end of the discussion before it has even really started. Apart from that, this doesn’t seem to offer any basis for a somehow convincing reaction to scepticism. Contextualism of this

kind would be much less attractive for those contextualists who also expect a reply to the sceptical challenge from contextualism. As mentioned above: The contextualist couldn't secure an ordinary context from sceptical threats anymore.

One might try to avoid all this and just say that what matters here is epistemological closeness or remoteness, not ontological or any other kinds of ranking. There is a certain plausibility to this: Aren't we talking about epistemology here?⁸ But there are further problems if one chooses this kind of reply. What determines epistemological closeness or remoteness? Intuitively it seems pretty clear that a world in which I am a brain in a vat is epistemically much more remote from what I take to be the actual world than a possible world which differs from the actual world (or what I take to be the actual world) only insofar as in that world I know certain facts about monkeys which I don't know in the actual world. Epistemological closeness or remoteness seems to depend on how much I know in the relevant worlds and what the difference as to the amount of my knowledge is in these worlds: I know some things, it seems, in the actual world, a bit more in the close possible world (where I know all that stuff about monkeys) but I know nothing about the external world in the vat-world. "Knowing quite a bit" is much closer to "Knowing a bit more than that" than to "Knowing nothing". One might want to add considerations of "quality" to such ideas of quantity: It also matters how important the relevant pieces of knowledge are. So, it would be a weighed sum of the quantity and quality of knowledge that determines the epistemological closeness or remoteness of possible worlds.

⁸ It is interesting to see that Lewis chooses ontological criteria of remoteness; whether that is compatible with standard contextualism (given the remarks above) must be left open here (cf. Lewis 1973, 75-76, Lewis 1979, 472).

There are a lot of problems with this: How to weigh the quantitative against the qualitative aspect? What determines both aspects anyway? That is, how can one spell out in a reasonable way all that talk about importance of pieces of knowledge? And how could one possibly quantify how much someone knows? I don't want to go into any of these problems. Whether one can deal with them or not doesn't concern me here. I only want to offer a rough picture of this version of the standard story and point at a much more simple problem with it: the threat of a vicious regress. According to our proposal, the truth-value of "Jack knows that he has hands" depends on how far out, epistemologically, the possibilities are that he has to eliminate (according to the attributor). How far out a given possible world is, depends on the difference as to what the subject knows in both worlds. In short, the truth-value of "Jack knows that he has hands" depends on what Jack knows in the actual world and on what he knows in all other relevant possible worlds. Since we are contextualists here, we should - when we are taking things strictly - stay at the meta-linguistic level and say this: The truth-value of "Jack knows that he has hands" depends on the truth-values of all sentences attributing knowledge to Jack as evaluated as true or false in the actual world and on the truth-values of all such sentences as evaluated as true or false in all the other relevant possible worlds. The truth-value of "Jack knows that he has hands" thus depends on the truth-value of many, many other knowledge sentences. It seems obvious that the truth-value of those other sentences will, again, depend on the truth-value of still other sentences. Hence we have an infinite regress (the alternatives of circularity or an arbitrary stop somewhere are not attractive). This is bad not so much because the explanans uses the explanandum (we don't have to try to define the concept of knowledge) but rather because it makes truth-values of knowledge sentences dependent

on each other and thus “hanging up” in the air. We could no longer understand what the truth conditions of some particular knowledge-sentences are because we’re always referred to the truth conditions of still other knowledge-sentences. One might, finally, try to explain epistemological closeness or remoteness of worlds not in terms of “knowledge” but rather in terms of other epistemological concepts, like justification. I won’t go into this but just remark that I have doubts that this will help a lot. Even if one can avoid the regress I would guess that the account would be too weak to explain ‘knowledge’ (see the problem mentioned in a similar context above).

The upshot of all this is that standard-contextualism à la DeRose and others faces serious difficulties. More generally, it seems very hard to imagine how one could come up with an interesting hierarchy of contexts that is in itself neither context dependent nor arbitrary. Sure, in some respects so-called "sceptical" contexts or standards are more remote than at least some of the so-called "ordinary" contexts or standards but there are other respects in which it is just the other way around. As Michael Williams has pointed out (1996; 2004), if a historian starts having doubts about the reality of the past she is not replacing laxer by stricter standards but rather changing the topic. Rather than a hierarchy of contexts and standards we should expect a diversity and plurality of them without a hierarchical order. If all this is true, then both the positive (standard) contextualist account of knowledge and the contextualist reply to the sceptic become quite implausible.

All this does, of course, not imply that the notion of epistemic standards should not play an important role for contextualists. It just cannot be used for the tasks some contextualists want to use it for. But there certainly are different standards corresponding to different contexts. Mary is a meteorologist and it might be true to say

of her in a lay-context that she knows that it will rain later today but not true in a professional context of her weather forecast lab: She hasn't run the typical tests yet.

Apart from that, there is a wide variety of kinds of things that could be called "standards". In some contexts, a true knowledge ascription presupposes that the subject meets certain standards of justification, etc. In other contexts, however, not a lot more than true belief might be sufficient for a true knowledge ascription (cf. Ernst 2002). Sometimes knowledge is compatible with a fixed "epistemic position", that is, an epistemic position that cannot – within the confines of the situation at hand - get better or worse (e.g., epistemic lottery situations); sometimes knowledge is not compatible with a fixed epistemic position (cf. Baumann 2004).

So, what we can learn from all this is not that contextualists should completely forget about the idea of standards. It is rather that they should put much (much!) less theoretical weight on it. Those variants of contextualism that work with remoteness-rankings of possible worlds run into insurmountable problems. And: There are other, rather neglected, aspects of the context-sensitivity of knowledge ascriptions that have little or nothing to do with epistemic standards. To that and the more constructive part of this paper I turn now.

2. Descriptions

I don't think one will ever be able to define the concept of knowledge. But even so, there is a very useful (partial) explanation of "knowledge": Knowledge requires a true belief that has been acquired in a reliable way. Usually, "reliability" is taken in an externalist sense (cf. Goldman 1986, 103; Dretske 1981a, ch.4) but I want to take it in a

much broader sense, including internalist ways of belief acquisition (good reasoning while being aware of the rules of correct reasoning, for instance, is a reliable way of belief acquisition). Let us call the way the subject acquires her belief her “method” of belief acquisition. Again, I want to use that word in a very broad sense: It should not imply that the subject uses the method intentionally or consciously (perception would be an example here).

There is a problem that has caused a lot of headaches for reliabilists: the so-called generality problem (cf. Feldman 1985, Alston 1995, Conee and Feldman 1998: Pollock 1984; Bach 1985; Brandom 1998; Beebe 2004, Hudson 2004). There is usually more than one way to specify the method of belief acquisition. Joe just acquired the belief that there is a desk in front of him but what exactly was the method? Perception, seeing, seeing-with-his-glasses-on, looking-at-furniture-with-his-glasses-on-while-suffering-from-a-remarkable-sleep-deficit? Depending on the specification of the method we can get different degrees of reliability, and thus, even different truth-values for my claim that Joe know there is a desk in front of Joe. Since there is no way to choose among these different specifications, reliabilism turns out to be an empty and therefore useless theory.

So much for a short description of the generality problem. I think it is a specific version of a much more fundamental problem, namely the reference class problem. What is the probability that an individual *a* is *F* (that Jack will live to the age of 72 years)? It depends on the relevant reference class *a* belongs to (the class of smokers, city dwellers, bikers, smoking city bikers, etc). Different reference classes give us different probabilities and the problem is to pick the one relevant reference class. There

is widespread pessimism concerning the possibility of a solution of the reference class problem.

The generality problem is, I said, just one aspect of this much broader problem. Even if we hold the method fixed (assuming that we have solved the generality problem), there will still be a reference class problem for reliabilism (cf. Wallis 1994). It doesn't only concern the individuation of the method but also the determination of the reliability of a given method. Does that mean that the prospects for reliabilism are even gloomier than the generality problem suggested? No, I think there is a different and quite constructive lesson to be learnt, and contextualism delivers it; in a sense, a particular version of contextualism is the "solution" of the generality problem and related problems (cf. also Heller 1995 and Cohen 1988, 115 who shortly mentions this issue). But let me explain. I will focus on one phenomenon which one can call "the time-sensitivity of reliability" (there is a parallel spatial case but I won't go into that here).

Consider the following example. Jack is the only witness in some court-case. The question is whether asking Jack about what happened is a reliable method of finding out the truth. This depends, of course, on how reliable Jack is. Suppose that Jack is reliable during the whole week in which the trial takes place. This is compatible with the fact that Jack suffers from a single momentary blackout between 5 to 10 am and 10 am on Wednesday. Only Jack notices that he suffers from a blackout. He happens to give correct answers to questions during his blackout. A single short blackout during the whole week does not put his general reliability as a witness during the whole week into doubt. Who is reliable at every moment in time? Suppose further that Jack is being interrogated just between 5 to 10 and 10 on Wednesday. With regard to this time-span

of five minutes, Jack is, of course, not reliable. Suppose further that the trial takes place during summer and that Jack is a periodical drunkard who consumes alcohol only during the cold seasons (and then in abundance). If Jack is drunk, then he is completely unreliable. The situation then is as follows:

Jack

- is no reliable witness between 5 to 10 am and 10 am on Wednesday;
- he is a reliable witness during the whole week;
- he is no reliable witness with regard to the whole year.

Similar things are true, *mutatis mutandis*, for the reliability of finding something out by relying on Jack's testimony. We can *ad libitum* consider further time-intervals with regard to which Jack might or might not be a reliable witness.

The important general point is that reliability (e.g., of Jack's testimonies and of reliance on his testimonies) is or can be time-sensitive. There can be and often is a great variation of probability with chosen time spans (like in the example just given). In that case, there is no reliability *tout court* but only reliability with respect to some time or time interval (and unqualified talk about reliability loses its meaning). Reliability is time-sensitive if the corresponding probability distribution over time is not constantly yielding the same value (or very similar values) but rather markedly different values for different times or time intervals. A constant probability distribution might be the exception. Hence, we can assume that normally or very often probability (and thus, reliability) varies with time.

A classical answer to this kind of problem says that one should choose the smallest reference class as the relevant one (cf. Reichenbach 1994, 383; on different but ultimately similar lines cf. Hempel 1965, 53-79, 397-403, and Salmon 1966, 90-92; cf. also Beebe 2004, 181). The problem with this and other, similar, proposals is not just that sometimes there is no smallest reference class (is there a shortest time interval?). More seriously, even if we have identified a smallest reference class the class will often be too small to allow for statistical information. It might be very small or even just consist of the one element we are considering. Furthermore, it is very hard to even make sense of probabilistic statements concerning individual cases or very small samples (but this is, of course, not uncontroversial and leads to more general questions concerning the interpretation of probability).

There are more problems. Take our Jack-example. Why should the shortest time-interval be the relevant one in the first place? Let us consider the proposal that the 5 minutes during which Jack is indeed questioned by the judge constitute the relevant time span. Why should that be so? One might say “Well, because that is when Jack gave his answers!” But what does that have to do with Jack’s reliability? And, by the way, weren’t his answers all correct? The point here is a bit subtle. We are interested in the 5 minutes of the interrogation because we are interested in the content and truth-value of Jack’s answers but not because we are interested in Jack’s reliability as a witness. That is a different issue. Consider this analogy. My notebook is reliable (in a non-epistemic sense). But just now it happens to freeze. Does that mean it is not reliable? Or not reliable at this moment? No, it just means that even a reliable notebook freezes from time to time.

What would we say if we changed the example a bit and assumed that Jack was fine (no blackout or anything like that) during the time of questioning but never fine (say, always drunk or hungover) at other times? I don't think we would be forced to consider him reliable; many would rather want to say he is not reliable, even if his answers happened to be correct, and even if he was reliable during the questioning period.

Apart from that, one attributor might be quite "liberal" and just require that the witness be reliable during his testimony whereas another attributor might "raise" the bar and require that the witness ought to be reliable during the whole day, week, month, etc. There is no reason why the latter should be considered "wrong" in any sense.

Even if we assume that the time during which Jack is being questioned is the relevant one here: What determines his reliability at that time, between 5 to 10 and 10? In order to answer that question, we need to determine the relevant reference class, again: Is it the class of all summer interviews? Or of interviews done when his stomach feels in this particular way? We are getting into a regress. In other words, there seems no way to get rid of the indeterminacy of the relevant reference class. One might propose to determine the relevant reference class at the time between 5 to 10 and 10 by using a possible worlds approach: Jack's reliability during the relevant time depends on how easily he could have gotten things wrong. But what determines the closeness or remoteness of possible worlds in which Jack gives false answers? This leads back to the objections made above against the idea of a remoteness-ranking of possible worlds.

What does all this imply? It seems that there just is no single right or relevant temporal reference class. We, the attributors of knowledge and reliability, must pick one – whether we are clearly aware of this or not. In the end, it all depends on our practical and theoretical interests. The prosecution might have an interest to use Jack in

other trials as a witness, too; hence, it will (tend to) judge Jack's reliability according to a longer time-interval. The defence might be rather interested to interrogate Jack in this particular phase of the trial (at the beginning of the Wednesday session); it will thus (tend to) judge Jack's reliability according to a shorter time interval (e.g., the interval between 5 to 10 am and 10 am). Now, given what I have said above, different interests and different criteria of relevance yield different views about the corresponding relevant probabilities and, thus, different views of Jack's reliability and of the reliability of the judge's method of interrogating Jack. And all those different descriptions are equally legitimate and adequate. This does, of course, not exclude that we very often agree amongst each other in our practical interests or on the criteria of relevance. Sometimes, there are even established and institutionalized rules (e.g., for what counts as good evidence or as trustworthy witnesses in court). And sometimes, there are such rules because without them we would not agree with each other.

We can thus say the following, given that knowledge presupposes the reliability of the method used: The truth-value of "The judge knows (by Jack's testimony) that the car was parked at the crime scene" varies with the attributor's view or description of the relevant temporal framework. Analogous points can be made about space. And finally, there is also the initial generality problem concerning the individuation of the method used. So, with respect to all these three aspects we should, I think, give a contextualist analysis of knowledge sentences. In this sense, contextualism "solves" the generality problem and the underlying reference class problem for "knowledge": It doesn't make the point go away, it just tells us that it really isn't a big problem after all.

This brand of contextualism is, of course, different from what I have called "standard-contextualism": It doesn't have to do with standards but rather with

descriptions (of spatio-temporal frameworks and of methods of belief acquisition). One could call it “description-contextualism”. The latter form of contextualism does not seem to suffer from the problems of the former that I discussed above.⁹ However, I don’t want to give up all forms of standard-contextualism but rather reject only those forms of it that are close enough to DeRose’s and related ones. I think that all this also shows that contextualism comes in more than one variety. It has at least two dimensions: standards and descriptions (perhaps there are even more than these two). A context dependency in one dimension does not imply anything about context dependencies in the other dimension; in that sense both aspects are independent from each other. To cut a long story short: There is more than one legitimate version of contextualism.¹⁰

9 One more (potential) relative advantage: Description-contextualism is not or at least not to the same degree susceptible to “warranted assertability maneuvers“ (cf., e.g., DeRose 1999, 195-203). It is hard to see how one could “wam“ description-contextualism.

10 I would like to thank Martijn Blaauw, René van Woudenberg, an anonymous referee, and an audience at Amsterdam Free University for comments and discussions. For discussion and comments on former versions of this paper I am grateful to Berit Brogaard, Duncan Pritchard, and Barry Smith. I would also like to thank audiences in Göttingen, Tübingen, Erfurt, Santiago de Compostela, and Dublin.

References

- Alston, William P. 1995, How to Think about Reliabilism, in: *Philosophical Topics* 23, 1-29.
- Bach, Kent 1985, A Rationale for Reliabilism, in: *The Monist* 68, 246-263.
- Baumann, Peter 2004, Lotteries and Contexts, in: *Erkenntnis* 61, 415-428.
- Beebe, James R. 2004, The Generality Problem, Statistical Relevance and the Tri-Level Hypothesis, in: *Noûs* 38, 177-195.
- Brandom, Robert 1998, Insights and Blindspots of Reliabilism, in: *The Monist* 81, 371-392.
- Cohen, Stewart 1987, Knowledge, Context, and Social Standards, in: *Synthese* 73, 3-26.
- Cohen, Stewart 1988, How to Be a Fallibilist, in: *Philosophical Perspectives* 2, 91-123.
- Cohen, Stewart 1998, Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery, in: *Australasian Journal of Philosophy* 76, 289-306.
- Cohen, Stewart 1999, Contextualism, Skepticism, and the Structure of Reasons, in: *Philosophical Perspectives* 13, 57-89.
- Conee, Earl and Feldman, Richard 1998, The Generality Problem for Reliabilism, in: *Philosophical Studies* 89, 1-29.

- DeRose, Keith 1992, Contextualism and Knowledge Attributions, in: *Philosophy and Phenomenological Research* 52, 913-929.
- DeRose, Keith 1995, Solving the Skeptical Problem, in: *Philosophical Review* 104, 1-52.
- DeRose, Keith 1999, Contextualism: An Explanation and Defense, in: John Greco and Ernest Sosa (eds.), *The Blackwell Guide to Epistemology*, Oxford: Blackwell, 187-205.
- Dretske, Fred I. 1970, Epistemic Operators, in: *Journal of Philosophy* 67, 1007-1023.
- Dretske, Fred I. 1981a, *Knowledge and the Flow of Information*, Cambridge/MA: MIT Press.
- Dretske, Fred I. 1981b, The Pragmatic Dimension of Knowledge, in: *Philosophical Studies* 40, 363-378.
- Ernst, Gerhard 2002, *Das Problem des Wissens*, Paderborn: mentis.
- Feldman, Richard 1985, Reliability and Justification, in: *The Monist* 68, 159-174.
- Fine, Kit 1975, Critical Notice [of Lewis 1973], in: *Mind* 84, 451-458.
- Goldman, Alvin I. 1986, *Epistemology and Cognition*, Cambridge/MA: Harvard University Press.
- Goodman, Nelson 1972, Seven Strictures on Similarity, in: N. Goodman, *Problems and Projects*, New York: Bobbs-Merrill, 437-447.
- Hawthorne, John 2004, *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Heller, Mark 1995, The Simple Solution to the Problem of Generality, in: *Noûs* 29, 501-515.

- Heller, Mark 1999, The Proper Role for Contextualism in an Anti-Luck Epistemology, in: *Philosophical Perspectives* 13, 115-130.
- Hempel, Carl G. 1965, *Aspects of Scientific Explanation and other Essays in the Philosophy of Science*, New York/ London: Free Press/ Collier.
- Hudson, Robert G. 2004, The Generality Problem, in: *Southern Journal of Philosophy* 42, 193-211.
- Jackson, Frank 1977, A Causal Theory of Counterfactuals, in: *Australasian Journal of Philosophy* 55, 3-21.
- Lewis, David 1973, *Counterfactuals*, Oxford: Blackwell.
- Lewis, David 1979, Counterfactual Dependence and Time's Arrow, in: *Noûs* 13, 455-476.
- Lewis, David 1986, *On the Plurality of Worlds*, Oxford: Blackwell.
- Lewis, David 1996, Elusive Knowledge, in: *Australasian Journal of Philosophy* 74, 549-567.
- Neta, Ram 2003, Contextualism and the Problem of the External World, in: *Philosophy and Phenomenological Research* 66, 1-31.
- Pollock, John 1984, Reliability and Justified Belief, in: *Canadian Journal of Philosophy* 14, 103-114.
- Reichenbach, Hans 1994, *Wahrscheinlichkeitslehre – eine Untersuchung über die logischen und mathematischen Grundlagen der Wahrscheinlichkeitsrechnung*, in: Hans Reichenbach, *Gesammelte Werke*, Braunschweig etc.: Vieweg, vol 7.

- Salmon, Wesley C. 1966, *The Foundations of Scientific Inference*, Pittsburgh, PA: University of Pittsburgh Press.
- Slote, Michael A. 1978, Time in Counterfactuals, in: *Philosophical Review* 1987, 3-27.
- Tversky, Amos 1977, Features of Similarity, in: *Psychological Review* 84, 327-352.
- Unger, Peter 1986, The Cone Model of Knowledge, in: *Philosophical Topics* 14, 125-178.
- Wallis, Charles 1994, Truth-Ratios, Process, Task, and Knowledge, in: *Synthese* 98, 243-269.
- Williams, Michael 1996, *Unnatural Doubts. Epistemological Realism and the Basis of Scepticism*, Princeton: Princeton University Press.
- Williams, Michael 2004, Knowledge, Reflection and Sceptical Hypotheses, in: *Erkenntnis* 61, 315-343.