

Contextualism about Knowledge and Justification by Default

Marcus Willaschek, Frankfurt/M.¹

Contextualism about knowledge is the view that the standards someone must meet in order to know something vary with the context of ascription. In this paper, I want to defend and refine a contextualist approach to knowledge and scepticism. After a brief exposition of the sceptical problem, I will sketch the standard contextualist approach to it as expressed (with significant variations), for instance, in the work of David Lewis, Steward Cohen, and Keith DeRose, and argue that this approach is unconvincing, among other reasons because it is too hospitable to the sceptic (1). Looking at knowledge-attributions in real-life cases will motivate a contextualist approach enriched by a default-and-challenge conception of justification (2), as has been proposed before by, among others, Michael Williams. Although I sympathise with much of Williams' account, I will argue that his conception of a 'default justificational status' is insufficiently complex, which in turn makes his version of contextualism also overly hospitable to the sceptic (3). Next, I will sketch some features of a sufficiently complex contextualist-cum-default-and-challenge conception of knowledge. Perhaps the most remarkable feature of the contextualism defended here, as compared to other versions of contextualism, is that it does *not* imply that the same knowledge-claim can be true in one context but false in another. This in turn is a consequence of construing contextualism about knowledge not as a linguistic thesis about the usage of the expression 'to know' and its cognates, but rather as a claim about the different standards at work in different epistemic practices. As I will argue, if a knowledge-claim can be evaluated by the standards of different practices, it is always the 'strictest' practice that counts (4). Despite incorporating some features of 'absolutist' conceptions of knowledge, however, the 'epistemic practice' contextualism defended here can deal with the sceptical challenge in a satisfactory way (5).

1. I leave my apartment. In the staircase, I stop and ask myself whether I locked the door. Did I? Do I know that I did? In order to know that I did, I must be able to rule out that I forgot to lock the door, which on reflection I can't. So I don't know that I locked the door. – Now I turn back to check whether I locked the door. I press the handle and find the door is locked. Do I *now* know that the door is locked? The intuitive answer clearly is: yes, now I know the door is locked.

¹ This paper has appeared in *Grazer philosophische Studien* 74 (2007), 251-272.

But many epistemologists would hesitate. After all, there are many possible situations compatible with my finding the door locked (more precisely, with my believing the door is locked after pressing the handle) in which the door is not locked. Here are some such situations:

- (a) The handle might be stuck.
- (b) I might not have pressed hard enough.
- (c) Someone might have tricked me by holding the handle on the other side of the door.
- (d) I might only have dreamt to have checked the handle.
- (e) I might only have dreamt to have checked the handle because everything I seem to experience is really part of a dream.
- (f) There might be no door, no handle, no hands, since I am a deluded brain in a vat.²

And, one could argue, as long as I can't rule out (a) - (f), I don't know that the door is locked.

Now possibilities (a) - (f) clearly fall into two distinct categories: while (a) - (d) can be empirically checked, and thus it is possible to rule them out at least in principle, (e) and (f) can never be ruled out by empirical investigation. (e) and (f) are so-called 'sceptical possibilities': If they obtain, there is no evidence that they do obtain; if they don't obtain, there is no evidence that they don't obtain, because all possible evidence is equally compatible with both their obtaining and their not obtaining.

Many philosophers think that sceptical arguments of the following type have a strong intuitive plausibility (the 'argument from ignorance')³:

- (1) I don't know that I am not a brain in a vat.
- (2) If I don't know that I am not a brain a vat, then I don't know that I locked the door.
- (3) I don't know that I locked the door.

Since the same argument can be run for any proposition about the so-called external world, it seems to follow that we really don't know most of what we take ourselves to know.

² Michael Williams offers a structurally similar list (concerning travel guides) in Williams 2003.

³ DeRose 1995.

One account of why this argument seems so compelling is offered by the sceptic: the argument seems to be compelling because it is correct. It is deductively valid and its premises and its conclusion are true. We really cannot know such things as whether the door is closed, whether I have two hands, and so on. However, almost everyone thinks that scepticism is false. So we have to explain not only why the sceptical argument seems compelling, but also where it goes wrong. Most of the various kinds of contextualism that have been proposed over the last 25 years aim primarily at meeting this double explanatory demand.

One prominent version of contextualism, developed in different ways by David Lewis, Keith DeRose, and Steward Cohen,⁴ is the view that has been called ‘semantic’ contextualism.⁵ According to this position, the context to be considered in evaluating knowledge-claims and knowledge-attributions is the conversation in which the claim is put forward or the attribution is made. On this view, there are two basic kinds of contexts: ordinary and sceptical. Ordinary contexts are characterized by low epistemic standards, which means that in order for *S* to know that *p* only some error-possibilities must be ruled out, namely those which are considered to be relevant in the context of a particular conversation in which knowledge is ascribed.⁶ (An error-possibility for the belief that *p* is any possibility in which non-*p*.) In a sceptical context, by contrast, epistemic standards are high: in order for *S* to know that *p* all alternatives to *p* must be ruled out. Most conversations generate an ordinary context, but once sceptical scenarios like the brain in a vat scenario are posited, a sceptical context is created; the epistemic standards change from low to high, with the result that all ordinary knowledge-claims turn out to be false in the context of that conversation. Note that this means that while two subjects *S*₁ and *S*₂ may possess the same evidence with respect to their beliefs that *p*, *S*₁ may know, and *S*₂ may not know, that *p*, if their respective contexts differ. Even more paradoxically, this means that a statement of the form ‘*S* knows that *p* at time *t*’ may be true in one conversational context, but false in another, even if everything about the subject, the available evidence, and the facts of the matter is constant.

In ordinary contexts, in which sceptical hypotheses are not considered, we may easily know such things as that we locked the door. All we have to do is, for instance, remember that we locked the door, or check or double check if we feel uncertain. But once sceptical hypotheses are considered, we enter a different context in which remembering or checking is

⁴ Cf. e.g. Cohen 1988, DeRose 1995, Lewis 1996.

⁵ Cf. Prichard 2002.

⁶ This formulation is meant to be neutral with respect to internalist (Cohen) and externalist (Lewis, DeRose) versions of semantic contextualism.

not good enough. In effect, in this sceptical context nothing is good enough. Therefore, in a sceptical context, we do not and cannot know that we locked the door (or know anything else about so-called external reality).

The intuitively compelling character of the sceptical argument is thus explained in the very same way it is explained by the sceptic: the argument is compelling because it is valid and its conclusion is correct. On the other hand, however, full-blown scepticism is avoided because the victory of the sceptic is limited to sceptical contexts: in ordinary contexts, where sceptical hypotheses are not at issue, we continue to know what we ordinarily take ourselves to know, including the fact that we locked the door. But even though scepticism is avoided, according to semantic contextualism the sceptic wins every argument simply by mentioning sceptical scenarios. I now want to argue that this can't be right.

2. Let us return to the list of possibilities supposedly incompatible with my knowing that I locked the door. As mentioned before, these possibilities fall into two classes: those that can at least in principle be empirically ruled out and those that can never be ruled out. The latter, of course, are the ones that give rise to sceptical problems. Now imagine again that I have just left my apartment, but that this time I am not alone, but with a friend. Halfway down the stairs the friend asks me whether I locked the door. Let us consider three different cases:

Case a: I answer 'Yes', and that is the end of it. Since an unqualified assertion typically expresses a claim to knowledge,⁷ by answering 'Yes' I claim to know that I locked the door. If this claim goes unchallenged, no further argument or reason-giving is needed. (Whether this means that no evidence is needed at all is a question to which I will return.)

Case b: Again, my friend asks whether I locked the door, but this time I am not sure and return to check. Again I find the door is locked. But my friend is not satisfied and suggests: 'The handle might be stuck'. How would I react? It depends. If the handle had been stuck before, I might say: 'Yes, you're right, I'd better check the handle'. In that case, I would perhaps unlock the door with my key, try the handle when the door is open, etc. But what if the handle never got stuck before, or if I fixed it only recently? Then the rational answer to the suggestions that the handle might be stuck would clearly be: 'No, don't worry, the handle's okay. The door is closed all right'.

⁷ Williamson 2000, 11f.; cf. Williams 2001.

Case c: My friend asks whether I locked the door, I say yes, she asks how I know, I respond that I clearly remember that I locked the door only a minute ago, and now my friend objects: ‘But perhaps you only dreamt that you locked the door’. What would I say to this? Probably my first reaction would be one of puzzlement: ‘What do you mean? How can you think that I only dreamt I locked the door when you were with me when we left the apartment?’ Now again there are two possibilities: Either my friend has a good answer to this question; for instance, she might point out that I am severely sleep-deprived and apt to fall into second-long periods of sleep without my noticing. In this case I might return to check the door, thereby admitting that I do not really know that the door is closed. But what if my friend has no such answer to give? What if the only answer comes to this: ‘It is logically possible and compatible with all available evidence, including our apparent memories to the contrary, that you only dreamt that you locked the door’. Would I, should I, return and check the door? The obvious answer is no. I would neither return, nor would it be rational for me to do so. Rather, I would insist that I did lock the door, thereby claiming to know that I did. And the same obviously goes for the two sceptical scenarios that I might be dreaming all the time and that we all might be deluded brains in vats: if this is all my friend had to offer in order to question my claim to knowing that I locked the door, the only rational reaction would be to point out that this may be possible, but that these far-fetched possibilities are irrelevant unless backed by further considerations, and to insist without any further argument that I know that I closed the door.

So we find that, at least in ordinary, non-philosophical contexts, there is a distinction between error-possibilities we do (and ought to) take seriously and others we don’t (and needn’t). What is missing in the cases in which we don’t take an alternative seriously is what Peirce once called a ‘positive reason’⁸ for doubting the belief in question. If I am sleep-deprived, or if, as in the earlier situation, the handle is old and tends to get stuck, these are positive reasons for doubting. Put differently, in order to question my belief that I locked the door, what is needed is *evidence to the contrary*. But a mere logical possibility is no evidence at all, neither for nor against something. It is logically possible that unbeknownst to me I just inherited a fortune from an uncle whom I have never heard of before. But the fact that this is logically possible is no evidence at all for the belief that I just inherited a fortune. There is only one

⁸ Peirce 1868, 140.

way in which logical possibilities can be evidence for a belief about something actual—namely, by ruling out that it is logically impossible that the belief is true. Apart from that, the mere fact that it is logically possible that p can never serve as evidence for the belief that p — and, as the examples just considered show, it can also not serve as evidence against the belief that not- p .

But further lessons can be drawn from the examples just considered.

First, our real-life practice of putting forward, questioning, and justifying knowledge claims follows what Robert Brandom and others have called a *default-and-challenge* pattern⁹: in many situations, knowledge-claims are considered to be legitimate without any backing by explicit justification or reason-giving. However, this default status is open to challenges.

Second, for a knowledge-claim to enjoy default status, it must meet certain conditions. Here contextual factors come into play. For instance, knowledge-claims that agree with general and/or expert opinion will typically enjoy default status, while heretical claims require explicit justification. Also, a claim about whether or not p issued in situations that are conducive to correct judgement about p will typically be default, while claims about something that is out of the subject's standard range of epistemic access will need explicit backing. If I claim that I just locked the door to my apartment, then this claim will typically enjoy default status, while my claim that Angela Merkel just locked the door to her office in Berlin will not. The reason is that memory affords a standard means of epistemic access to what *I* just did, but not to what someone else just did far away from me.

What is remarkable about this is that the deliverances of memory do not play the role of evidence: for my claim that I did lock the door to enjoy default status it is not necessary that I really remember that I locked the door. What counts is that generally people *do* remember correctly — and thus may be *supposed* to remember correctly—what they did a minute ago. Whether I really remember or not becomes relevant only when the reliability of my memory is being questioned. (The same holds for other modes of knowledge-acquisition such as perception and testimony.) I will return to the question of preconditions for enjoying default status in more detail below.

The third and final lesson I want to draw is this: not any old logical possibility can serve as a challenge to a knowledge-claim that enjoys default status. *It is simply not true that the mere mentioning of error-possibilities undermines a claim to knowledge.* This means that the semantic contextualist diagnosis of scepticism can't be correct, because it assumes that mentioning or thinking of sceptical scenarios suffices to raise the epistemic standards in such

⁹ Brandom 1994, 177, Williams 2001, 25.

a way that we no longer know what we do know in ordinary contexts. What is needed is more than that, namely evidence to the contrary: either evidence that the claim is in fact false or evidence that the conditions required for default status are not fulfilled. For instance, the fact that there is no non-circular argument for the reliability of memory as a source of information about the past is not enough to challenge my claim that I just locked the door. Rather, what is needed is evidence that *my* memory, in this very situation, is not reliable. Again, contextual factors come into play, since the distinction between relevant and irrelevant challenges, challenges that must be answered and those that can reasonably be shrugged off, depends on the given situation and specific circumstances: For someone suffering from a loss of short-term memory, the question of whether one really remembers to have locked the door poses a serious challenge; to someone whose memory works properly, it does not.

Notice that this is not just externalist reliabilism: What depends on the context (e.g. on the reliability of someone's memory) is not simply the question of whether the person knows that *p* or not, but rather what standards the person must meet in order to know that *p*: is it enough for the person in the given situation to correctly believe that *p* (knowledge by default), or must she adduce further evidence and rule out contrary evidence?

3. A version of contextualism that has taken to heart the first lesson about the default—and-challenge structure is the one developed by Michael Williams, most clearly in his 2001 book *Problems of Knowledge*.¹⁰ Williams keeps to the traditional conception of knowledge as justified true belief, but interprets it within a default-and-challenge model of justification, in which both default status and the relevance of challenges depend on various contextual factors.

Williams distinguishes between two kinds of justification: personal and evidential.¹¹ Personal justification concerns the question of whether the epistemic subject acts epistemically responsibly in believing and claiming what she does. Evidential justification, by contrast, consists in the 'adequate grounding' of a belief by the available evidence. Both personal justification and objective well-groundedness are required for a belief to count as knowledge. The default-and-challenge model primarily concerns the subject's personal justification: I am personally justified in believing that *p* either if my belief that *p* enjoys default status or if I can counter all relevant challenges (including standing objections to *p*). As standard externalist scenarios such as 'barn façade county'¹² show, personal justification

¹⁰ Williams 2001.

¹¹ Williams 2001, 22.

¹² Cf. Goldman 1976.

of a true belief is not enough for knowledge: if I correctly believe that there is a barn over there, my belief may count as being justified by default, but if what I see is the only real barn among many façades, I still do not know that I see a barn. So Williams additionally requires what he calls the ‘adequate grounding’ of a belief, which depends in part on the external circumstances in which the belief is formed and held (Williams 2001, 162).

Williams’ version of contextualism differs from semantic contextualism in yet another important respect. Whereas the latter view traces all changes in epistemic standards to moves in a conversation, Williams takes into account a broad variety of contextual factors, which he brings under five types¹³: (i) *semantic* (What must be presupposed for the question of whether *p* or not *p* to make sense/to arise?); (ii) *methodological* (What must be presupposed for our methods of inquiry to work; e.g. the reality of the past with respect to methods of historical research?); (iii) *dialectical* (What challenges are actually being advanced? Which are generally considered to be standing objections?); (iv) *economic* (How important is the correctness of the knowledge-claim to the subject and/or the attributor? How much effort is necessary to rule out certain error-possibilities?); and (v) *situational* (Is the chosen method of belief-formation in fact reliable in the given situation?). Since all these factors can influence the epistemic standards according to which a given knowledge-claim must be evaluated, there is no clear-cut distinction between high and low standards, but rather a complex net of standards, any of which can be strict in one respect but rather loose in others.

Obviously, the default-and-challenge model of justification plays a central role in Williams’ complex version of contextualism. But what about the other two lessons that can be drawn from the examples I considered?

The second lesson was that in order to enjoy default status a belief has to fulfil certain conditions, such as being orthodox, being formed under standard conditions for the reliable working of the chosen method of belief formation, and so on. By contrast, Williams holds that ‘one is entitled to a belief or assertion [...] in the absence of appropriate “defeaters”’¹⁴. So a belief is justified by default, according to Williams, as long as it is not faced with serious challenges. Of course, one might argue that the heterodox character of a belief, or the fact that it was formed under non-standard conditions, should count as an ‘appropriate defeater’. But this seems to me to mislocate their epistemological impact: The fact that in Copernicus’ time almost everyone believed that the sun moves around the earth is no ‘appropriate defeater’ to Copernicus’ claim that the earth moves around the sun; if it were, Copernicus would have had to convert the majority of his contemporaries to a heliocentric view in order to know that the

¹³ Cf. Williams 2001, especially ch. 14.

¹⁴ Williams 2001, 149.

earth moves around the sun. The heterodox character of Copernicus' claim simply means that his claim did not enjoy default status and that, in order for his claim to count as knowledge, Copernicus had to *argue* for his view. Among other things, he had to give appropriate answers to the substantial objections raised by his contemporaries. One such objection, for instance, was that one ought to expect the movement of the earth to produce an airstream; but there is no such airstream (an objection answered convincingly only by Galilei). The fact that most people *believed* that the sun moves around the earth, by contrast, was no such substantial objection.¹⁵ What this shows is that being heterodox rules out default status not in virtue of being a challenge or defeater, but rather because orthodoxy of a belief is a prerequisite of its enjoying default status in the first place.

The same seems to be true about the proper working of the method of belief-formation one employs. Imagine I look out of the window in plain daylight and correctly claim that it's snowing. Someone watches me and asks 'How do you know?' Assuming the person saw that I looked out of the window, I wouldn't know how to respond except by saying something the person already knows, namely that I know it is snowing because I looked out of the window. And it seems that in the absence of contrary evidence I am not epistemically required to give any further explanation of how I know that it is snowing. – But now imagine that I sit in a room with the windows closed. Last time we looked outside it wasn't snowing. All of a sudden I claim that it's snowing. You ask: 'How do you know?' or 'How can you tell?' In this case, I am obviously required to give an informative explanation in order to be counted as knowing that it is snowing. It would not suffice to appeal to something the other person already knows, such as: 'Well, I have been sitting in the room all the time; so of course I know that it just started snowing'. What is called for is further information about my method of belief-formation. For instance, I might respond that I have a scar that itches every time it starts snowing. If I can't produce a convincing answer, I must withdraw my claim to know that it is snowing. Since the same bare question is being asked in both cases ('How do you know?'), and since in both cases no evidence against my claim is put forward, it is hard to see why the question should be a serious *challenge* in the second case but not in the first. The difference rather seems to be that in the second case, but not in the first, the question is meant to remind me that in the given situation my claim does not enjoy default status, but must be backed by explicit reason-giving.

¹⁵ Obviously, there is no clear-cut distinction between objections that are and those that are not substantial. In general, substantive objections tend to focus on the specific content of the claim in question, while considerations that undermine a claim's default status tend to concern the subject's epistemic and dialectical position. But there may be exceptions.

Hence we must distinguish between two ways in which a belief can fail to enjoy default status: (1) by not meeting the requisite conditions such as being orthodox or being formed under standard conditions, and (2) by being faced with serious content-specific challenges.

I now turn to the third lesson I mentioned, namely that challenges to claims that enjoy default status need to be backed by evidence against the challenged claim, and that mere logical possibilities do not constitute evidence. Williams is aware of this: ‘A defeater does not come into play simply in virtue of being mentioned: there has to be some reason to think that it might obtain’.¹⁶ On the other hand, however, he claims in the same book that challenges based on sceptical hypotheses, once put forward, ‘deprive ordinary knowledge-claims of their default justificational status’,¹⁷ and even endorses the view that ‘we may temporarily lose our knowledge when we project ourselves into the rarified context of “doing epistemology”’.¹⁸ But a sceptical hypothesis is such that there *can* be no evidence either for or against it, since its obtaining and its not obtaining are equally compatible with all possible evidence. So with respect to a sceptical hypothesis, there is never a ‘reason to think that it might obtain’. It therefore seems to me that Williams is not consistent in claiming on the one hand that challenges must be based on evidence (‘reasons to think that they obtain’) and conceding on the other that sceptical hypotheses ‘deprive ordinary knowledge-claims of their default justificational status’. We should stick with the former claim and reject the latter.¹⁹

4. Let us pause for some reflections about the general characteristics of the contextualism that emerges out of the critical discussion of semantic contextualism on the one hand and Williams’ version of contextualism on the other. First, the resulting kind of contextualism is not primarily a linguistic thesis about the correct use of the expression ‘to know’ and its cognates. What depends on the context is the correctness of attributions of knowledge and justification, but these attributions typically take a non-linguistic form. If I hail a taxi, for instance, I attribute to the driver the knowledge that my behaviour means that I need a taxi. One might object that all I need to attribute to the driver is the corresponding belief. But if the default-and-challenge account sketched above is correct, then the driver’s belief will typically be justified by default; and since the belief in question is true, there is no reason not to regard it as knowledge. If I were to give linguistic expression to what I attribute to the driver, it

¹⁶ Williams 2001, 161; cf. 150f.

¹⁷ Williams 2001, 186.

¹⁸ Williams 2001, 195; also cf. Williams’ discussion of the ‘instability of knowledge’ in Williams 1991. In a similar way, David Lewis has claimed that knowledge is ‘elusive’: We have it only as long as we don’t reflect on it.

¹⁹ But cf. Williams 2003, 990f., where Williams himself criticises contextualists such as Lewis and DeRose for being too hospitable to scepticism.

would be most natural to say that, being a taxi driver, he *knows* what my hailing a taxi means. So much for non-linguistic attributions of knowledge. But even attributions that take a linguistic form will typically not employ the expression ‘to know’ and its cognates. If I ask you what time it is and you answer ‘It’s two o’clock’, you implicitly claim to *know* that it’s two o’clock. Otherwise, you ought to have said something like: ‘I’m not sure, but I believe it’s two o’clock’. As mentioned before, straightforward assertions that *p* typically have the force of self-attributions of knowledge that *p*. And if I react to your assertion by going to the train station in order to catch a train at half past two, then I accept your knowledge-claim and thus attribute knowledge to you. What depends on the context is the correctness of knowledge-attributions in this sense (and thus the conditions for knowledge). Only a small fraction of knowledge-attributions take the form of sentences containing the expression ‘to know’. Thus, the contextualism I want to defend is not primarily a linguistic thesis. As I want to show now, it also does not imply that the truth-value of knowledge-attributions, linguistic or not, *varies* with the context.

In order to see why this is not the case, it is necessary to say a bit more about what a context is. We have already rejected the idea that the kind of context relevant for knowledge-attributions is simply the context of a conversation. I now want to suggest that the context on which the correctness of knowledge-claims depends is constituted by two factors: an epistemic practice and the relevant facts of the matter. Epistemic practices are rule-governed social practices in which knowledge is acquired and attributed according to specific epistemic standards: sciences such as biology or physics, practices such as law, medicine, engineering, etc. Besides these highly regimented and self-reflective epistemic practices there is the vast variety of social practices such as crafts, commerce, and sports, each of which has some specific epistemic standards of its own, but which mostly employ the same all-purpose set of epistemic standards that governs commonsense attributions of knowledge. Let’s call this latter our ‘ordinary’ epistemic practice. The epistemic standards employed in different practices overlap, but there are also important differences. In the empirical sciences, for instance, knowledge is tied to the possibility of empirical confirmation; in mathematics and related formal disciplines, knowledge requires proof; in the law, knowledge from testimony is restricted by certain formal procedures such as taking an oath; in various crafts, practitioners can tell things apart simply by looking or touching, while laypersons can do so only by indirect methods; etc. etc. These standards determine (a) what kinds of beliefs enjoy default status; (b) what kinds of error-possibilities are relevant challenges; (c) what counts as answering a challenge (as ruling out an error-possibility); and (d) what counts as establishing

a claim that does not enjoy default status. However, these standards are *general* in the sense that they must be applicable to a broad variety of cases. In order to determine the epistemic status of a *specific* belief in a given situation, something else must be taken into account--namely, the relevant facts of the matter.

By this I mean those facts on which it depends *which* error-possibilities satisfy the criteria of relevance of a given epistemic practice in the given situation. For instance, before there were fake Rolex watches on the market, one could know there was a Rolex watch before one by simply reading the brand name on the watch face; the mere logical possibility that someone might counterfeit Rolex watches was no relevant challenge as long as there was no indication that someone really did so. However, if there are thousands of fake Rolex watches on the market, the possibility that the watch before me is a fake becomes relevant – even if I don't know about the existence of fake watches. If I cannot tell the fake from the real thing and if there is a significant number of fakes around, then I simply don't know that this watch before me is a Rolex (assuming that all I have to rely on is my own judgement). I suggest that we analyse this situation as follows: There is a general standard implicit in our ordinary epistemic practice according to which one knows that *p* only if one can rule out all *relevant* error-possibilities. Obviously, the specific criteria of relevance will vary with the epistemic practice and the facts of the matter. Here are some examples of criteria of relevance implicit in our ordinary epistemic practice: A possibility's being relevant requires *some* reason to think that it in fact might hold; merely logical possibilities are never relevant in this sense; for an error-possibility to be relevant, its obtaining must be consistent with generally accepted knowledge as well as with specific knowledge about the subject's situation; an error-possibility is relevant if the kind of error in question is common or to be expected under the circumstances; etc. etc. Now the facts of the matter determine, with respect to the case at hand, *which* error-possibilities satisfy the criteria of relevance specified by the epistemic standards: in the situation we imagined, the possibility that the watch is a fake is relevant since there are many fakes out there, and thus there is some reason to think that the watch might be a fake. By contrast, the possibility that one is only hallucinating a watch is not relevant since there is no reason to think that one is hallucinating. In the imagined situation, then, one knows that one holds a watch in one's hand, but one does not know that it's a Rolex. The epistemic standard (rule out all error-possibilities for which there is reason to think that they may in fact obtain) and the facts of the matter (many fakes out there) taken together thus determine whether an error-possibility is relevant or not. In this sense, knowledge *depends* on context.

I will call this kind of contextualism ‘epistemic practice contextualism’, because it is epistemic practices that set the standards for knowledge and justification. However, this is not to deny that a specific context also includes the facts of the matter that determine which error-possibilities are relevant in the given case. So changes in the relevant facts constitute changes in context if the factual changes affect the relevance of error-possibilities – even if the epistemic practice remains the same.

That knowledge depends on context, however, does not necessarily mean that the truth of individual knowledge-claims *varies* with the context, so that the same claim can be true in one context and false in another. The reason for this is that the correctness of a knowledge-attribution in *any* context typically depends on the standards of the *strictest* context in which the claim can be evaluated. Let’s call this latter context ‘the standard-setting context’. In the case of Rolex watches, for instance, the standards are set by the context of watch-making: if an expert watchmaker opens the watch, disassembles the clockwork and finds that the watch is a genuine Rolex product, then the possibility that the watch is a fake has been effectively eliminated. No further confirmation is needed, and none is possible. We cannot increase our confidence, for instance, by handing the watch over to a physicist and asking him to check whether it’s genuine. Even though the epistemic standards employed by the physicist may be in some respects ‘stricter’ than those of the watchmaker, this typically won’t help in determining whether the watch is genuine or not, since physicists, as such, are not experts in distinguishing genuine from counterfeit watches.²⁰ This suggests that the epistemic standards with respect to a given knowledge-claim about some subject matter *M* are set by the ‘expert practice’, if there is any--that is, the practice of those people who know best about *M* and to whom one defers judgement in that matter. Whether someone knows that *p* thus depends on whether the assertion that *p* is justified according to the epistemic standards of the ‘expert practice’. Otherwise, the layperson, simply because her standards are less exacting, could know things about a scientific subject matter that experts do not know, which seems absurd. (Of course, the difference between experts and laypersons does not only concern the epistemic standards, but also, and even more importantly, the ability to determine the facts of the matter.)

However, a layperson need not become an expert in order to know what the experts know, since implicit in our ordinary epistemic practice is the standard that one is justified in one’s beliefs about a subject matter *M* if one bases one’s beliefs about *M* on what is

²⁰ Of course, there may be exceptions – perhaps a fake can be distinguished from the real thing only on a molecular level. What this shows is that the assignment of knowledge-claims to epistemic practices, too, depends on the relevant empirical facts of the matter.

considered as knowledge about M by the respective experts. The expert justification of the belief need not be available to the layperson; still, her being justified in her belief about genuine Rolex watches, supernovae, or the human genome ultimately depends on the belief's meeting the epistemic standards of the expert practice. Such are the benefits of the epistemic division of labour.

Note, by the way, that sceptics, by definition, are not experts in any field. By admitting that the standards are set by the 'strictest' practice in which a knowledge-claim can be evaluated, the epistemic practice contextualist does not surrender to the sceptic, since even the strictest epistemic practices distinguish between relevant and irrelevant error-possibilities and allow for default justification. As we've already seen, even the watchmaker typically need not be able to rule out that he is dreaming in order to know that a watch is a genuine Rolex. I will return to the issue of scepticism in the last section.

For many subject matters there are no experts – or, equivalently, everyone's an expert: It doesn't take a meteorologist to determine whether it's raining here and now; all it takes is average perceptual capacities. Here, the standard-setting context is just our ordinary commonsense practice of evaluating knowledge-claims. But even in this case, changing to the scientific context of meteorology does not undermine the validity of our ordinary epistemic standards for the question of whether one knows that it's raining here and now: if someone knows that it's raining here and now according to ordinary standards, it is hard to see how a meteorologist might prove her wrong (according to meteorological standards). The reason is that with respect to the question of whether it's raining here and now, meteorologists rely on the same epistemic standards as ordinary folk.

So the situation is this: either there is an expert practice with respect to subject matter M , in which case someone has knowledge about M only if her true belief is justified according to the standards of that practice. Or there is no expert practice, in which case someone has knowledge if her true belief is justified according to our ordinary epistemic standards. Either way, changing the context will not change the truth value of a knowledge-claim. Hence, the truth-value of knowledge-claims *depends* on the context, but doesn't *vary* with the context. So in one important respect epistemic practice contextualism is invariantist. Nevertheless, it is a kind of contextualism in the sense of the opening sentence of this paper-in that different epistemic practices constitute different contexts with different epistemic standards. Relativism, however, is avoided by recourse to the epistemic division of labour: If there are no experts, our ordinary standards count; if there are experts, however, their standards kick in. Obviously, there can be much controversy about who's an expert, what the relevant standards

are, and whether or not a knowledge-claim meets the relevant standards. But these questions and the controversies about them make sense only if there can be correct and incorrect answers to them, which presupposes a non-relativist epistemology.²¹

But what then about an observation that seems to speak in favour of semantic versions of contextualism, namely that we tend to attribute knowledge rather generously in some contexts, but more parsimoniously in others? For instance, if Peter tells Mary today that he will go to Munich tomorrow, then it seems absolutely appropriate for Mary to say tomorrow that she knows that Peter is in Munich. (Imagine Peter and Mary work in the same office in Frankfurt and someone asks where Peter is and Mary answers: ‘I know where he is! He’s in Munich. He told me himself’. There seems to be nothing wrong with what Mary says.) But now imagine that later Peter is accused of a crime in Frankfurt that took place on the very day he was allegedly in Munich. It seems that Mary can no longer say that she knows where Peter was if all she has to rely on is his announcement that he was going to go to Munich. The reason is that now error-possibilities have become relevant that were not relevant before. So it seems that the correctness of a knowledge-claim can vary with the context after all.

In response, the epistemic practice contextualist can make use of what Keith DeRose has called a ‘Warranted Assertibility Maneuver’²²: The weaker standards according to which Mary knows that Peter was in Munich do not concern the *truth-conditions* of her knowledge-claim, but only the conditions of *warranted assertibility* of knowledge. In other words, it was appropriate for Mary to *say* that she knew, even though in fact she didn’t know. In many situations it may be appropriate for pragmatic reasons to ascribe knowledge without making recourse to the ‘strictest’ standards. For instance, if nothing much hinges on whether Mary is right or not, or if checking is difficult, then her belief will pass as knowledge. But it seems that when we ask whether she ‘really’ knows, pragmatic and economical considerations just don’t count. Perhaps the possibility that Peter did not go to Munich originally *seemed* irrelevant to Mary, but then it turned out to be relevant after all. So what we should say is that, even though it became apparent only afterwards, Mary never knew that Peter was in Munich. (Again, this does not imply surrender to the sceptic as long as we hold fast to the idea that even the strictest standards for a given knowledge-claim require a distinction between relevant and irrelevant error-possibilities.)

²¹ Therefore, the kind of contextualism I want to defend does not imply that explicit knowledge-ascriptions change their truth-value with changes in context. This means that the linguistic arguments against contextualism (e.g. Schiffer 1996; Stanley 2004; Capellen/Lepore 2005) do not apply.

²² DeRose 1995; of course, DeRose discusses ‘warranted assertibility maneuvers’ as moves made by the critics of (semantic) contextualism. The epistemic practice contextualist can acknowledge the limited force of these since her view is not primarily about assertions, but about knowledge.

Much more needs to be said in order to defend epistemic practice contextualism against its contextualist and non-contextualist rivals. Here my aim is only to give a rough sketch of a version of contextualism that promises to avoid the problems and weaknesses of both the standard semantic contextualisms and Williams' non-semantic version of contextualism. Working out the details of epistemic practice contextualism would primarily require a detailed account of the criteria of relevance employed in various epistemic practices, a project which would be of interest quite independently of most epistemologists' obsession with scepticism. Having said this, I will now return to the question of how to respond to the sceptical argument.

5. What, if anything, is wrong with the 'argument from ignorance'? Williams' diagnosis of scepticism is centred around the idea that the sceptic must presuppose what Williams calls 'epistemological realism'²³: the view that our beliefs and claims fall into natural epistemological kinds, like 'beliefs based on perception,' 'beliefs based on memory,' and so on, and that the beliefs stand in fixed evidential relations by virtue of the epistemic kind to which they belong. Only epistemological realism, so Williams claims, allows the sceptic to question all our knowledge-claims, or all our knowledge-claims of a certain kind, at once. Otherwise, the sceptic would have to proceed in a piecemeal fashion, investigating the epistemic standing of one individual claim after another, which of course means that the sceptic wouldn't get anywhere near the general conclusion that we don't know anything, or at least that we don't know anything in various broad domains.

This is an illuminating diagnosis, but it is difficult to see how it can be brought to bear on the sceptical argument considered above:

- (1) I don't know that I am not a brain in a vat.
- (2) If I don't know that I am not a brain a vat, then I don't know that I locked the door.
- (3) I don't know that I locked the door.

Of course, this argument, if successful, only shows that I don't know that I locked the door. But no theory such as epistemological realism is needed in order to generalise this result; all we need is the class of propositions incompatible with my being a brain in a vat. Which propositions belong in that class is a matter of logic and linguistic meaning. And if we have a

²³ Williams 2001, 170f., 191ff.; cf. Williams 1991.

grasp on the class of propositions incompatible with my being a brain in a vat, we can generalise the argument simply by pointing out that, if successful at all, it works for all propositions of that class and shows that they cannot be the content of legitimate knowledge claims.

So a different diagnosis is called for, and I think that the default-and-challenge model of justification affords such a diagnosis. To see this, consider a situation in which we put forward the first premise of the sceptical argument not as part of a syllogism, but as a challenge in a conversation: A taxi-driver claims to know how to get from here to Goethestrasse; you point out that this can't be right because he doesn't even know that he is not dreaming.²⁴ Assuming the taxi-driver is not, as it might happen, a former student of philosophy, and assuming further that he considers you worthy of an answer at all, he will not, and ought not, withdraw his knowledge-claim, but will rather insist that he knows the way alright and also knows that he's not dreaming. If you are a very insistent person, you may go on and explain to the taxi-driver why he can't possibly know that he's not dreaming: namely, because of the fact that any evidence for his being awake might in principle be a part of his dream. But if the taxi-driver has his wits about him, he will respond that he remembers how he woke up in the morning after a good night's sleep, how he went to work, and how you boarded his taxi and began a silly conversation. True, he might say, all this is compatible with the possibility of my being asleep and only dreaming this up, but this possibility in itself is no good reason to doubt my well-established belief that I'm awake. The mere fact that it is conceivable that right now he is dreaming does not go against his being awake, as long as there is no positive evidence in favour of the possibility that he is dreaming.

Of course, it is highly improbable that you hit on a taxi driver who will give you a speech like that. But I think that this speech spells out what many people think when they are confronted with sceptical arguments: they simply don't see their point, and insist that they do know the sceptical hypotheses to be false. This is the first step of my diagnosis of the sceptical argument: I deny that it even *seems* to be compelling to non-philosophers. And the reason is not that non-philosophers 'don't get it', but rather that they rightly and rationally insist on the principle that challenges to knowledge-claims – knowledge-claims that enjoy default status – require evidence against the correctness of the knowledge-claim. The mere possibility of an error, unaccompanied by any indication that it may actually have occurred, is

²⁴ I change the sceptical scenario from 'brains in a vat' to 'dreaming' merely to increase the plausibility of the narrative; challenging knowledge-claims in ordinary real-life situations by recourse to brain-in-a-vat scenarios would most probably evoke reactions of complete puzzlement – and questions about the mental sanity of the challenger.

no such evidence. So the first step of my diagnosis, if correct, achieves two things: It denies the first explanandum assumed by most philosophers, namely that the sceptical argument seems to be compelling; and at the same time explains the second explanandum, namely that the sceptical argument goes wrong somewhere. It goes wrong in the first premise, that we don't know the sceptical hypotheses not to obtain. Of course we cannot argue, in G. E. Moore-like fashion, from our having hands to our not being brains in a vat. But this is because having hands is no *independent* evidence for not being a brain in a vat, so that we can't argue from the one to the other.²⁵ But still we can know that we are not brains in a vat, because the belief that we are not enjoys default status, and no evidence against it has been offered. In fact, it is part of the idea of sceptical hypotheses that no evidence in their favour is even possible.

The second step of my diagnosis then must concern the question why the sceptical argument seems so compelling to many *philosophers*; in particular, to explain why many philosophers find it plausible that we do not know the sceptical hypotheses not to obtain. The answer I want to propose is that in philosophy we abstract from the real-life contexts in which knowledge-claims are being issued and evaluated.²⁶ We don't ask ourselves whether this or that person in this or that situation knows the way to Goethestrasse, but whether anyone can ever know the way to Goethestrasse, or, for that matter, whether anyone can know not to be a brain in a vat. What we ask is whether, and how, anybody can know anything at all.²⁷ The effect of this de-contextualizing tendency of philosophical reflection is that there seems to be no distinction between the relevant and the irrelevant, between reasonable and unreasonable challenges, because these distinctions are highly context-sensitive. Once we abstract from context, all challenges can seem equally relevant. This is why I think it is important not just to consider sceptical syllogisms, but to imagine real-life situations in which knowledge-claims are being issued and criticised. This helps us keep in mind that there are error-possibilities which, because they are not backed by any evidence, are simply irrelevant for the question of whether we know the way to Goethestrasse. If we forget this feature of our epistemic practice, however, we will come to the conclusion that we cannot know the sceptical hypotheses not to obtain, since after all there is the uneliminated and ineliminable possibility that they might in fact obtain. Only if we keep in mind that being uneliminated is not enough for an error-

²⁵ Cf. Wright 2002.

²⁶ Cf. Putnam 1998.

²⁷ Cf. Stroud 1989.

possibility to be relevant can we see that we know very well that we are not brains in a vat (provided that we are not brains in a vat).²⁸

²⁸ I would like to thank Hannes Ole Matthiessen and Andreas Maier for valuable discussions about the topic of this paper and for extremely helpful critique of various earlier versions of it.

References

- Brandom, Robert 1994: *Making It Explicit*, Cambridge (Mass.)
- Cappellen, Herman/LePore, Ernest 2005: *Insensitive semantics: a defense of semantic minimalism and speech act pluralism*, Malden (Mass.): Blackwell
- Cohen, Steward 1988: "How to be a Fallibilist", *Philosophical Perspectives* 2, 91-123
- DeRose, Keith 1995: "Solving the Sceptical Problem", *Philosophical Review* 104, 1-52
- Goldman, Alvin I. 1976: "Discrimination and Perceptual Knowledge", *Journal of Philosophy* 73, 771-791
- Lewis, David 1996: "Elusive Knowledge", in *Australasian Journal of Philosophy* 74, 549-567
- Peirce, Charles S. 1886: "Some Consequences from Four Incapacities", *Journal of Speculative Philosophy* 2, 140-157
- Prichard, Duncan 2002: "Two Forms of Epistemological Contextualism", in *Grazer Philosophische Studien* 64, 19-55
- Putnam, Hilary 1998: "Scepticism", in M. Stamm (ed.), *Philosophie in systematischer Absicht*, Stuttgart: Klett-Cotta 1998, 239-268
- Schiffer, Stephen 1996: "Contextualist Solutions to Scepticism", in *Proceedings of the Aristotelian Society* 94, 317-333
- Stanley, Jason 2004: "On the Linguistic Basis for Contextualism", in *Philosophical Studies* 119, 119-146
- Stroud, Barry 1989: "Understanding Human Knowledge in General", in M. Clay/K. Lehrer (eds.), *Knowledge and Skepticism*, Boulder, 31-49
- Williams, Michael 1991: *Unnatural Doubts. Epistemological Realism and the Basis of Scepticism*, Oxford
- Williams, Michael 2001: *Problems of Knowledge*, Oxford
- Williams, Michael 2003: "Skeptizismus und der Kontext der Philosophie", in *Deutsche Zeitschrift für Philosophie* 51/6, 973-991
- Williamson, Timothy 2000: *Knowledge and Its Limits*, Oxford
- Wright, Crispin 2002: "(Anti-)Sceptics Simple and Subtle: G.E. Moore and John McDowell", in: *Philosophy and Phenomenological Research* Vol. LXV No.2, 330-348