

ELKE BRENDEL and CHRISTOPH JÄGER

CONTEXTUALIST APPROACHES TO EPISTEMOLOGY: PROBLEMS AND PROSPECTS

ABSTRACT. In this paper we survey some main arguments for and against epistemological contextualism. We distinguish and discuss various kinds of contextualism, such as *attributor contextualism* (the most influential version of which is *semantic, conversational, or radical contextualism*); *indexicalism*; *proto-contextualism*; *Wittgensteinian contextualism*; *subject, inferential, or issue contextualism*; *epistemic contextualism*; and *virtue contextualism*. Starting with a sketch of Dretske's Relevant Alternatives Theory and Nozick's Tracking Account of Knowledge, we reconstruct the history of various forms of contextualism and the ways contextualists try to handle some notorious epistemological quandaries, especially skepticism and the lottery paradox. Then we outline the most important problems that contextualist theories face, and give overviews of their criticisms and defenses as developed in this issue.

1. INTRODUCTION

Contextualist approaches to epistemological concepts and problems have become extremely popular in contemporary epistemology. "Contextualism," however, is just an umbrella term for a wide variety of theories. Their common starting point is the thesis that the truth values of knowledge ascriptions (or ascriptions of epistemic justification) are *context-dependent*. This context-dependency is said to provide the key to resolving some of the most notorious epistemological quandaries, including the skeptical problem and the lottery paradox. In working out this idea, contextualist approaches begin to diverge. One major family of views has come to be called *attributor contextualism*, the most influential form of which is *semantic or conversational contextualism*. This approach has most prominently been advocated by Stewart Cohen, David Lewis, and Keith DeRose. The other major strain is *subject contextualism*, one of the leading proponents of which is Michael Williams. In what follows, we shall sketch some main steps in the historical development of conversational and subject contextualism and outline the core characteristics and philosophical targets of these positions. Second, we will outline some crucial problems and objections contextualist accounts face,



Erkenntnis 61: 143–172, 2004.

© 2004 Kluwer Academic Publishers. Printed in the Netherlands.

and provide overviews of the defenses as well as the criticisms and alternative proposals presented in the papers in this issue.

2. RELEVANT ALTERNATIVES, TRUTH-TRACKING, AND EPISTEMIC CLOSURE

One theory of knowledge which has had a major impact on recent contextualist approaches is the so-called “Relevant Alternatives Account” first proposed by Fred Dretske in the early 1970s and further developed by Gail Stine and others.¹ According to Dretske, an epistemic subject *S* knows that *p* (at time *t*) only if *S* is in an epistemic position that allows her to eliminate all *relevant alternatives* to *p* (at *t*). A proposition *q* is an *alternative* to *p* just in case *q* entails not-*p*. Yet, according to Dretske it is not necessary, in order to know *p*, that one be able to exclude *all* the alternatives to *p*. What is required instead is merely the ability to eliminate or rule out certain *relevant* alternatives. So what makes an alternative relevant? This depends on the epistemic situation. Usually during an ordinary visit to the zoo, the possibility that the animals you take to be zebras are cleverly disguised mules is an irrelevant alternative, and it is thus not necessary that you be able to rule it out in order to know that the animals are zebras. But now suppose, for example, that it is well-known that the zoo director, in order to save money, often disguises common animals as exotic animals and occasionally puts cleverly disguised mules in the zebra paddock. In this situation, the mule alternative becomes relevant and, it would seem, you do *not know* that the animals you are looking at are zebras, unless you can rule out the possibility that they are cleverly disguised mules – even if they are in fact zebras (cf. Dretske, 1970, p. 1016).²

Dretske’s painted-mule scenario is a situation where *local* or *restricted skepticism* – in this case skepticism about whether in some particular situation appearances are trustworthy – threatens some knowledge claim. But his account is also designed to provide an answer to *global* or *radical skeptical arguments* such as the notorious brain-in-a-vat (BIV) argument:

BIV:

- (1) I don’t know that I am not a (handless) brain in a vat.
- (2) If I don’t know that I am not a (handless) brain in a vat, then I don’t know that I have hands.
- (3) I don’t know that I have hands.

The *skeptical paradox* consists in the fact that such arguments are valid and use premises that intuitively seem true. Yet we are not willing to accept the conclusions. Dretske's original answer is, very roughly, that skeptical scenarios are *irrelevant* alternatives. If so, premise (2) in the above argument turns out to be false.

Add to this view that what counts as a relevant alternative is determined by the alternatives that are salient for the person ascribing the epistemic attitude (or lack of it) to the subject, and the position you arrive at is attributer contextualism. In his contribution to this issue, Dretske explicitly distances himself from such forms of contextualism (which he calls *radical contextualism*). Because of the deep influence his theory has had on such forms of contextualism, however, his view may well be called a kind of *proto-contextualism*.³ We shall come back to Dretske's position below.

Another highly influential account of knowledge is Robert Nozick's "tracking" analysis.⁴ The question of whether *S* knows that *p* in a given situation depends, according to Nozick, not only on *S*'s having a true belief that *p*, but also on certain counterfactual relations between *p* and *S*'s believing that *p*. In particular, two subjunctive conditionals must be satisfied: (1) If *p* had been false, *S* would not have believed that *p*; i.e., *S* knows that *p* only if, in the nearest possible worlds in which *p* is false, *S* no longer believes that *p*; and (2) if *p* were true, then *S* would have believed that *p*, i.e., in all the closest worlds where *p* is true, *S* believes that *p*. Given these conditions, *S* can know that she has hands, even though *S does not know* that she is not a brain in a vat: One of the nearest possible worlds in which *S* does not have hands is a world in which *S*, for example, lost her hands in an accident; and in this world she would not believe that she has hands. Furthermore, in all the closest worlds in which it is true that *S* has hands, *S* believes that she has hands. Since *S*'s belief that she has hands satisfies both truth-tracking conditions, it follows from Nozick's account that *S* knows that she has hands. However, *S does not know* that she is not a brain in a vat, since in the nearest possible world in which *S is* a brain in a vat (let us assume that *S* is not one in the actual world), *S* would still believe that she is not a brain in a vat. Thus our ordinary knowledge claims, as in the relevant alternatives account, can still be true, even if we don't know that the skeptical hypotheses are false.

One consequence of Dretske's and Nozick's theories is the failure of what many regard as a highly plausible epistemic principle, namely, the *principle of epistemic closure* (PEC). According to PEC, knowledge is closed under known logical entailment. PEC can be roughly stated as follows:

PEC: If S knows that p and knows that p implies q , then S also knows that q .⁵

The skeptic implicitly appeals to PEC when she argues as follows: Since we don't know that we are not brains in vats, and since we clearly know that having hands implies not being a (handless) brain in a vat, it follows that we don't know that we have hands. Since similar reasoning can be applied to any other proposition about some ordinary fact, the skeptic concludes that we don't have any knowledge of such facts.

Now, as we have already seen, Nozick's account entails that we *can* know that we have hands without knowing that we are not brains in vats, although we know that having hands implies not being a brain in a vat. Rejecting PEC also allows Dretske to avoid radical skeptical conclusions. In his classic papers on the topic, Dretske argues that PEC holds only when the entailed proposition's negation is a relevant alternative to the proposition in question. Furthermore, since, at least in everyday situations, being a brain in a vat is not a relevant alternative to having hands, we need not know that we are not brains in vats in order to know that we have hands – even though we clearly know that having hands implies not being a (handless) brain in a vat. In recent work, Dretske puts forward the view that, even in contexts where skeptical alternatives *are relevant*, rejecting an unrestricted closure principle is the appropriate response to the skeptic. With regard to “heavy-weight” implications such as the negation of skeptical hypotheses, he maintains that closure does not hold even when such hypotheses have become salient (cf. Dretske, 2004).

Whether or not we want to regard skeptical hypotheses as relevant alternatives, the main problem with Dretske's and Nozick's accounts is that rejecting PEC is a high price for solving the skeptical problem. PEC is after all a highly plausible principle of knowledge acquisition. DeRose contends that not knowing that we are handless brains in vats, while at the same time knowing that we have hands, is an “abominable conjunction” and an “intuitively bizarre result” (DeRose, 1995, p. 201). Conversational contextualism, championed by Cohen, Lewis, and DeRose, attempts to solve the skeptical problem by appealing to the context-sensitivity of knowledge claims without giving up closure.⁶

3. CONVERSATIONAL CONTEXTUALISM

The main claim of conversational contextualism (henceforth: CC) is that the sentence “ S knows that p (at t)” can be true in one

conversational context and false in another – for the same subject *S* and the same proposition *p* (and the same time *t*).⁷ According to CC, it is always the context of the *speaker* that determines the truth conditions for a given utterance of “*S* knows that *p*.” CC is therefore a version of attributor contextualism. Except in cases of self-ascriptions of knowledge, the subject’s conversational context plays no role in determining the truth conditions for “*S* knows that *p*.”

But how does a knowledge-ascriber’s context determine the semantic standards of a knowledge claim, and what induces context changes? According to CC, the raising and lowering of standards is determined by *conversational features*. If the speaker’s attention is drawn to an error-possibility that has not yet been considered for a proposition *p*, the standards are then raised. In particular, this means that the mere *mentioning* of some error-possibility *e* makes *e* salient which, in turn, causes a shift from a lower-standards context (where *e* need not be ruled out) to a higher-standards context (where *e* must be ruled out). Even if an epistemic subject *S* meets the standards for knowledge put in place by a low-standards context where an error-possibility *e* to *p* is not salient, “*S* knows that *p*” may still turn out to be *false* in a higher-standards context where *e* is salient, if *S* cannot rule out *e*. As we will see, a number of contributors to this issue object to the idea that context changes are solely induced by conversational features.

One main goal of CC is to give a satisfying response to the skeptical challenge while nevertheless explaining skepticism’s intuitive appeal. To put it in terms of relevant alternatives: Since in *everyday situations* the skeptical possibility that we are brains in a vat is an irrelevant alternative to our belief that we have hands, we don’t have to rule out this skeptical hypothesis in order to know that we have hands. But there might be situations where skeptical possibilities are relevant alternatives to our ordinary knowledge claims – for example, in the context of a philosophy seminar on epistemology. Since we cannot rule out the possibility that we are brains in vats, we cannot know *in these situations* that we have hands. So, on the one hand, the skeptical challenge is met, because our ordinary knowledge claims remain true as long as we are in a context of everyday life. On the other hand, the appeal of our skeptical intuitions is explained, since in philosophical contexts where skeptical possibilities are relevant, our ordinary knowledge claims turn out to be false.

The contextualist thesis can also be described by saying that knowledge claims are *indexical*. This, it is contended, provides a semantic explanation for the apparent fact that sentences of the form “*S* knows that *p*” can have different truth values in different

contexts. According to CC, the semantics of other indexical expressions – like “flat” – can serve as a model for understanding the indexicality of knowledge claims. Whether the assertion “ X is flat” is true depends on the standards of flatness determined by the context of utterance. Similarly, whether or not the ascription “ S knows that p ” is true depends on the epistemic standards put in place by the knowledge-ascriber’s context. But the contextualist’s indexicality thesis has been challenged. For example, in his contribution, Wayne Davis contends that there is compelling linguistic evidence against the indexicality of knowledge claims. The contextualist must also explain why competent speakers who can identify assertions that are uncontroversially indexical find it difficult to recognize the presumed indexicality of knowledge claims.⁸

Another main goal of CC is to provide a solution to the *lottery paradox*. Let us assume that S bought a ticket in a fair lottery and that the chances of this ticket winning are very low – 1:10,000,000. If S is the lucky winner, she will get 10 million dollars. Although there is overwhelming statistical evidence for the belief that S ’s ticket will lose, many people share the intuition that S nevertheless does not know that her ticket will lose. Let us assume furthermore that, given S ’s meager income and her lack of rich relatives, S claims to know that she will never be a multi-millionaire. Now we have a problem: S ’s knowing that she will never be a multi-millionaire seems to imply her knowing that she will not win the lottery – which contradicts the intuition that S fails to know that she will lose.⁹

Cohen, in particular, maintains that CC provides a solution to this version of the lottery paradox on the grounds that CC explains the widespread intuition that S does not know that she will lose:¹⁰ In ordinary-standards contexts the sentence “ S knows that she will never be a multi-millionaire” is true, and so is the sentence “ S knows that her ticket will lose.” But once we think about the lottery and the chance (however slight) her ticket has of being drawn, this remote possibility becomes salient and creates a context in which the standards for knowledge are so high that “ S knows that she will lose the lottery” is false. Thus according to CC there only *seems* to be a paradox, because of an unnoticed context change from one knowledge claim to the other. But in the case of the lottery paradox as well, the contextualist solution has been attacked. For example, Peter Baumann and John Greco challenge the assumption that the salience of chances of error can explain the intuition that S does not know that her ticket will lose. We shall outline their criticisms as well as Cohen’s reply below.

4. OBJECTIONS TO CONVERSATIONAL CONTEXTUALISM

One of the most general objections to CC is that, since it is a theory about the semantics of knowledge ascriptions, it is an exercise in the philosophy of language, rather than an epistemological position that provides insight into the nature of knowledge. This “meta-linguistic ascent objection” has been put forward and developed at some length by Ernest Sosa (2000).¹¹ DeRose, however, notes that to the extent that contextualism engages in the philosophy of language, it undeniably deals with issues that are of utmost importance to epistemology (1999, p. 188).

Another very popular objection to CC has already been touched upon: It seems counterintuitive to maintain that, simply by mentioning skeptical hypotheses (or drawing attention to them in some other way), we can deprive a person of her everyday knowledge. Moreover, critics have argued, it is just not true that people simply withdraw or object to knowledge claims when they are confronted with skeptical hypotheses. As Richard Feldman has illustrated (Feldman, 1999, p. 100): Suppose you are at a cocktail party and participate in a debate about the healthiest diet. Some people offer arguments for the view that it is healthy to eat lots of carbohydrates, others argue in favor of protein. After a while you chime in with the remark: “But at least I know this: I’m no brain in a vat!” According to CC, this assertion should provoke dissent, for in the contextualist’s view the mere mention of the BIV hypothesis (even in claiming to know that it is false) raises the epistemic standards. Thus your claim should be greeted at least with considerable epistemic suspicion. But this is not what happens. You may produce strange looks, but outside the philosophy classroom you will hardly succeed in provoking dissent with knowledge claims to the effect that some outlandish skeptical scenario does not obtain. (We concede that to some extent people’s reactions may also depend on how many cocktails they have already consumed.)

Objections along these lines, which can generally be classified as objections regarding the *dynamics of context shifts*, come up in several papers in this issue, such as in those of Antonia Barke, Wayne Davis, Fred Dretske, Mylan Engel, and Frank Hofmann. For example, both Davis and Engel charge that contextualism predicts – falsely – that when elevated skeptical standards are in force we will find ourselves converting to skepticism. In fact we don’t. Indeed, no one reading the papers in this volume (we hope) will cease to believe that there is an external world, that she has hands,

and that she is not a brain in a vat. Engel also claims that skeptical arguments tend to lose their force once we become familiar with them. If this is true, it shows that there is another kind of epistemic dynamic which needs explaining and which contextualism cannot account for.

The contextualist has two main replies to such worries. First, as especially Cohen has argued (see for example Cohen, 1999), the contextualist can incorporate an *error theory* into his account. According to an error theory, competent speakers are often unaware of, or systematically misled by, the context-sensitivity of knowledge ascriptions. We shall return to this topic below. Another answer the contextualist can offer is that attacks from such quarters are anyway misconstruing his project. His project, he might say, is *not* primarily the description of what happens if actual epistemic subjects encounter skeptical hypotheses. Instead, the issue is one of *normative epistemology*, and thus his proposal should not be evaluated with regard to people's actual feelings and reactions toward skepticism. However, suppose there *is* agreement that some knowledge claim which in different circumstances is true has, under the pressure of skeptical arguments, become false. Then we are still left with the question of how those inflated standards can be lowered again. The only way of regaining knowledge would seem to be to ignore, or forget, the skeptical possibilities that have become salient. But it is unclear whether, and if so in which way, this could happen. How exactly can we find our way back into epistemic naïveté? Do we reenter low-standards contexts as soon as we leave the philosophy classroom? And are we dragged into high standards again once we return from our coffee break or the cocktail party?

Lewis (1979, 1996), DeRose (1995) and Cohen (1999) are well aware of this problem. Nevertheless, the objector claims, they underrate its critical potential and are confronted with a problem of the type “paradox of epistemic laziness.” According to Lewis (1996, p. 222), on the contextualist view, epistemology turns out to be an “investigation that destroys its own subject matter.” Indeed, but the reverse side of the coin is that, the more epistemically blind we are, the more we know. Ignoring skeptical arguments puts us in a better epistemic position than we would be in if we engaged in critical reflections about our everyday knowledge claims. Yet, as Hofmann notes in his paper: Should not knowledge, however exactly one may want to analyze it, at least be construed as an *achievement*?

Let us now look more specifically into the main topics of the papers in this volume. In “Externalism and Modest Contextualism,” Fred **Dretske** notes approvingly that CC is inspired by a relevant alternatives account of knowledge. However, he rejects the direction in which proponents of CC have been steering with his account. The contextualist is committed to the view that once skeptical hypotheses are mentioned, we have moved to a context in which it is true to say that a given subject *never knew* the ordinary propositions that conflict with the skeptical hypotheses. In general, Dretske argues, *S*’s knowledge could apparently only be sheltered by insulating *S* from ever thinking about skepticism. Yet, *if* skepticism is false, isn’t it false in the philosophy classroom as well as in the grocery store? The general worry here is that contextualism is conceding far too much to the skeptic. Why are skeptical arguments so appealing? In a way, the contextualist answer is straightforward: Skepticism is *true* for all of us who have thought about skeptical arguments. Worse than that, according to CC, skepticism is even true for all of us whose knowledge claims have been subjected to skeptical contemplations *by others*. This, one may feel, is not a good resolution of skeptical problems. Dretske then prescribes what he thinks can cure such maladies: rejecting closure.

Dretske links this point to his overall externalist account of knowledge, which is cashed out in terms of a detailed theory of information. However, in his paper “Skepticism, Information, and Closure: Dretske’s Theory of Knowledge,” Christoph **Jäger** argues that, at least with regard to ordinary empirical propositions and their antiskeptical consequences, Dretske’s information-based externalism is in fact incompatible with his rejection of closure. Two of the most central and most influential features of Dretske’s epistemology appear to be mutually exclusive. On Dretske’s theory, *S* knows of some perceptual object (or source of information) that it exemplifies a certain property, if and only if there is some signal which carries the relevant information and which, in virtue of carrying that information, causes *S*’s belief that the object has that property. Furthermore, a signal is said to carry the information that *p* only if the probability of *p*, given the signal, is 1. But then, Jäger shows, on Dretske’s theory the relation of a signal’s carrying the information that *p* is closed under logical entailment. Second, Jäger draws on an embellished version of the closure principle and suggests a causal interpretation of the epistemic basing relation for Dretske’s account. He then shows that, given these assumptions, Dretske is committed to the view that, with regard to the propositions in question, also *knowledge* is closed

under known entailment. If so, Dretske must either abandon his information theory of knowledge, or must himself embrace skepticism. Both alternatives would have far-reaching consequences for his epistemology. And in either case, Jäger concludes, Dretske's answer to skepticism, as it stands, *cannot* be regarded as a viable alternative to contextualism.

Mylan Engel's paper "What's Wrong With Contextualism, and a Noncontextualist Resolution of the Skeptical Paradox," comprises two parts. First, he reconstructs and criticizes CC.¹² Besides putting forth criticisms from directions that have already been outlined above, Engel reminds us of the contextualist commitment to *unspeakable and unthinkable knowledge*. This phenomenon has been noted by DeRose (1995, section 13f.) and is also discussed by Davis: CC contends that when low standards are operant we know that we are not brains in vats. According to the contextualist's view of context dynamics, however, we are unable to assert *or even think* the proposition that we know this. (As soon as we would entertain this thought, standards would rise, and we would be entertaining a falsehood.) Engel then proceeds to develop an alternative diagnosis of the appeal of skeptical arguments, the central idea of which is that the skeptic is exploiting an equivocation between *metaphysical* and *epistemic possibility*. Engel contends that the skeptic either supports the first premise of the BIV argument by appealing to the "possibility" of one's being a BIV or argues directly for the skeptical conclusion using the following *argument from possibility*:

AP:

- (1) It is possible that I am a handless BIV.
- (2) If it is possible that I am a handless BIV, then it is possible that I don't have hands.
- (3) If it is possible that I don't have hands, then I don't know that I have hands.
- (4) Hence, I don't know that I have hands.

A crucial problem with this argument, Engel contends, is that it fails to specify whether, and where, it employs concepts of metaphysical possibility or, alternatively, of epistemic possibility (e-possibility). Engel then introduces a notion of e-possibility according to which a proposition is *e-possible* for you iff, roughly, (i) you don't know that it is false and (ii) you could not come to know that it is false on the basis of propositions you know. He then distinguishes between a fallibilistic and an infallibilistic sense of "e-possibility" and argues that three possible readings of AP in terms of e-possibility can be

rejected without further ado, since they are either equivocating between the fallibilistic and the infallibilistic senses of e-possibility or, in the case of a consistent infallibilistic interpretation, can be dismissed as uninteresting. The only version of AP that deserves at least *prima facie* to be taken seriously, according to Engel, is an interpretation in terms of fallible e-possibility.

EAP2f:

- (1) It is fallibly e-possible that I am a handless BIV.
- (2) If (1), then it is fallibly e-possible that I don't have hands.
- (3) If it is fallibly e-possible that I don't have hands, then I don't fallibly know that I have hands.
- (4) Hence, I don't fallibly know that I have hands.

Given Engel's account of e-possibility, the crucial first premise of EAP2f can be restated roughly as follows:

- (1)* I don't fallibly know that I am not a handless BIV, and I could not come to know fallibly that I am not a handless BIV, strictly on the basis of propositions I know.

Engel contends that EAP2f simply begs the question against those who believe in the truth of ordinary knowledge claims, because in asserting EAP2f.1 (i.e. (1)*), the skeptic presupposes that we don't know that we have hands.

Gilbert **Scharifi** begins his commentary on Engel's paper by questioning the role of Engel's reconstruction of AP. Does it differ substantially from the good old skeptical BIV argument (with which Engel started his discussion)? Scharifi identifies two different interpretations of Engel's anti-skeptical argument, a "straightforward interpretation" and a "sophisticated interpretation." In the straightforward interpretation, Scharifi argues, it turns out that e-possibility plays no role, since in that case Engel's question-begging charge against the skeptic could just as well have been directed against the original BIV argument: "It does not make any difference to this argument whether the skeptic's first premise reads 'It's e-possible for me that *b* [I am a handless BIV]' or 'I don't know that $\sim b$ '". Scharifi then turns to the "sophisticated interpretation" in which the notion of e-possibility, as he concedes, does play an essential role. Yet, Scharifi argues, Engel at best "gets a standoff": Maybe the skeptic cannot legitimately claim that we fail to know ordinary empirical propositions, Scharifi maintains, but then we cannot claim to know them.

Scharifi then takes the question-begging challenge head-on by trying to find an alternative reason for EAP2f.1 that doesn't appeal to the conclusion of EAP2f (i.e., that we don't know fallibly that we have hands). It is here that Scharifi invokes the sensitivity requirement in support of EAP2f.1. If knowledge requires sensitivity, then since we are not sensitive to the truth where our BIV-status is concerned, we fail to know that we are not BIVs (and so being a BIV is an e-possibility, just as EAP2f.1 asserts). Scharifi contends that Engel may of course reject the assumption that sensitivity to the truth is a necessary condition for knowledge. But, Scharifi argues, if Engel rejects the sensitivity requirement, there will no longer be any need for Engel's anti-skeptical argument. Thus, Scharifi argues, if the sensitivity requirement holds, Engel's question-begging charge fails, since in that case the skeptic has an independent reason for the first premise of his argument. If on the other hand the sensitivity requirement can be dismissed, he concludes, Engel's argument becomes superfluous.

Contextualists often try to back up their position with examples designed to illustrate that true knowledge claims become false when error possibilities are introduced. In "How to Be an Anti-Skeptic and a Noncontextualist" Bruce **Russell** takes a closer look at such examples and argues that there is a better explanation for the alleged intuition that people lose their knowledge when local or global skeptical hypotheses become salient. Consider DeRose's bank case: On a Friday afternoon Hannah and her husband are deliberating whether to deposit their paychecks at the bank. Noticing long lines at the counters, Hannah proposes returning the next day, saying: "I know the bank will be open tomorrow, since I was there on a Saturday morning just a few weeks ago." But then her husband points out to her that they must pay an important bill by Monday and that the bank may have changed its business hours. Hannah withdraws her claim and concedes that she does not really know that the bank will be open tomorrow. There is one possible interpretation of such examples according to which the subject simply stops believing the proposition once error possibilities have been introduced (and thus also ceases to know the proposition). DeRose, however, stipulates that in his bank example the relevant belief is not lost. Russell argues that, contrary to what the contextualist says, the best explanation of such cases (and of corresponding examples that work with radical skeptical hypotheses) is that we fallaciously infer that knowledge goes by the board. In fact, he argues, the subject still knows the proposition in question. Our implicit inference, he suggests, is this: "If S

knows that the bank will be open on Saturday, it must be open on Saturday. But it might not be open. So *S* fails to know that it will be open.” Yet this inference relies on a well known misinterpretation of the fact that what is known “cannot be false.” Nothing can be known unless it is true. But of course, this is only to be understood in terms of the *de-dicto* modality: Necessarily, if *S* knows that *p*, then *p*; it is not to be interpreted in terms of a *de-re* modality regarding *p*, i.e. as claiming that, if *S* knows that *p*, then necessarily *p*. Russell maintains that it is easy to confuse these modalities and that this explains why we are easily pulled in by such arguments. Moreover, he contends that we are epistemically blameless and thus (deontologically) justified in committing this fallacy.

Regarding cases in which people concede, when confronted with global skeptical arguments, that they don't know some ordinary empirical proposition they thought they knew, Russell endorses Engel's proposal and suggests that they are misled by subtle equivocations regarding “possible” and “knows.” But again, since these equivocations are subtle, Russell believes that people are deontologically justified when they are attracted by skeptical arguments. He claims that almost all arguments for skepticism fit this pattern. Against Scharifi, Russell utilizes an argument from Alvin Plantinga to show that sensitivity to truth is not a necessary condition of knowledge in any context. Russell concludes with one word of advice to anti-skeptical noncontextualists and with another to contextualists. To the noncontextualist: Whenever a contextualist believes that only his account can explain why we think that knowledge is lost under the pressure of skeptical arguments, the noncontextualist should try to explain why that argument merely looks good and why we are therefore blameless in accepting its conclusion, even if it rests on equivocations. Russell's advice to the contextualist is to avoid examples that start with either moderate or radical skeptical hypotheses, and to guard himself from fallacious skeptical inferences. Furthermore, a subject's justification should, according to Russell, be tied to the evidence that she *should* have, which may indeed be a matter of contextually determined interests. Finally, Russell recommends that contextualists dissociate themselves from all forms of skepticism. No one, he concludes, should be led to deny invariantism because of skepticism.

In his paper “Are Knowledge Claims Indexical?”, Wayne **Davis** focuses on the fact that proponents of CC explicitly present their account as a semantic theory about knowledge ascriptions. One curious feature of this semantics, i.e. the unassertability of true

knowledge ascriptions, has already been sketched. This, Davis notes, seems to be at least as abominable as the conjunctions licensed by a Dretske–Nozick account. Davis then reconstructs and criticizes CC in particular with respect to the indexicality thesis and proposes a pragmatic alternative. One of his main points is that there are clear and important asymmetries between indexical sentences and knowledge ascriptions. He also discusses in detail why he thinks the alleged contextualist solution *fails* to explain the skeptical paradox. The paradox arises in virtue of the following “inconsistent triad” (IT) of seemingly true propositions:

IT:

- (1) I know that I have hands.
- (2) If I know that I have hands, then I know that I am not a BIV.
- (3) I don’t know that I am not a BIV.

According to CC, context determines which of these three propositions we must reject: In ordinary-standards contexts where we are not contemplating skeptical scenarios, (1) and (2) are true, while (3) is false. In high-standards contexts where skeptical hypotheses are salient, (1) is false, while (2) and (3) are true. Thus, according to CC, the supposed paradox simply dissolves once we attend to the context of ascription. But, as Davis points out, most people think that the above inconsistent triad generates a fundamental paradox, since even after contemplating these three propositions, they all still seem true.

In reply to such objections, Cohen has put forth the *inattention thesis* (1988) and the *error theory* (1999): Our reluctance to accept or endorse the implications of elevated standards may simply be due to an inattention or blindness to context shifts. But, Davis asks, how then can Cohen maintain that knowledge claims are indexical? After all, competent speakers normally have no problem in correctly relating indexical utterances to the relevant contextual parameters. (When you say “I am cold,” I don’t normally mistake your utterance as referring to me.)¹³ Furthermore, ordinary indexical expressions typically don’t generate problems that are similar to the skeptical paradox. Davis goes on to present (what he regards as) more semantic evidence against indexicalism and then outlines an alternative pragmatic theory. Its basic idea is that the contextualist’s *explanandum*, i.e. the contextual variability of knowledge claims and ascriptions, may better, and at lower costs, be explained by invoking a Gricean theory of conversational implicature. Ordinary knowledge claims, even if they are false when evaluated under strict standards,

can nevertheless be conversationally appropriate. For they are often intended to be interpreted loosely, as only meaning or conversationally implying that we are, for present purposes, *sufficiently close to knowing*.

In his comment on Davis's paper, Gerhard **Ernst** makes three major points. First, he agrees with Davis that indexicalism does not provide an adequate answer to skeptical paradoxes, but he offers an independent reason in support of this verdict. Ernst contends that there is an alternative explanation of our reluctance to give up ordinary knowledge claims when we are confronted with skeptical arguments. This explanation is that we simply cannot, or don't *want* to, give up our deep-rooted ordinary beliefs, in much the way that a father, even when presented with overwhelming evidence that his son is a second Jack the Ripper, will desperately insist that he knows otherwise. (Descartes for instance was very clear about this phenomenon; it was precisely in view of this problem that he introduced the evil demon hypothesis in the First Meditation.) Ernst's second point is that Davis does not spell out his pragmatic alternative adequately. Moreover, he seems to favor an invariantist analysis of knowledge. Yet, Ernst worries, the history of epistemology of the last four decades strongly suggests that the prospects for invariantism are dim. Finally, Ernst critically examines Davis's asymmetry arguments. One of Davis's observations, for example, is that inferences of the form "*S* spoke truly when saying 'Smith knows that *p*,' therefore Smith knows that *p*" appear to be valid, whereas analogous inferences that involve indexicals ("*S* spoke truly when saying 'I am cold,' therefore I am cold") are clearly not valid. However, Ernst alleges, the indexicalist will answer that there is indeed no simple fact of the matter as to whether Smith knows or does not know. From the fact that it was true from one point of view to say "Smith knows that the flight stops in Chicago" it does not follow that this was also true from another point of view and thus that Smith also knows this in another context.

Another important question for conversational contextualists is what happens to the truth conditions of knowledge claims if the "conversational score" of the participants of a conversation is pushed in different directions. DeRose has recently raised this question and put forth his own solution – the so-called "gap" view (DeRose, forthcoming). In her paper "Keeping the Conversational Score," Verena **Gottschling** takes up DeRose's discussion of this problem. She first identifies two intuitions which seem to be plausible from a contextualist point of view: *the intuition of persisting individual*

standards, namely, the intuition that when participants of a conversation don't cooperate and instead insist on different standards for knowledge, their knowledge claims can have different truth-conditions due to these different individual standards; and the *intuition of contradiction*, which maintains that in cases of conflicting standards the participants contradict each other in ascribing different truth-values to a knowledge claim.

Gottschling then outlines different options available for a contextualist to determine the conversational score in cases of conflicting individual standards. (In order to simplify matters, she only considers conversations with two participants: a skeptic and a "commonsensian".) The "different scoreboard" view claims that each speaker has her own personal scoreboard and that there is no shared scoreboard. This view, so Gottschling argues, violates the intuition of contradiction. The "no scoreboard" view, in which knowledge claims lack truth-values in cases of conflicting individual standards, also violates the intuition of contradiction: Since these knowledge ascriptions lack truth-values, they don't contradict each other. On the other hand, "single scoreboard" views, in which there is only one scoreboard that determines the truth-conditions of a knowledge claim (either the score set by the skeptic, or the score set by the "commonsensian", or a "balanced scoreboard"), cannot accommodate the intuition of persisting individual standards.

After presenting these options, Gottschling turns to DeRose's new "gap" view according to which "*S* knows that *p*" is true if *S* meets the personally indicated standards of both speakers; it is false if *S* fails to meet either set of standards; and "*S* knows that *p*" (as well as "*S* does not know that *p*") is neither true nor false if *S* meets one set of standards but fails to meet the other one. She then argues that DeRose's view fares no better than the other views. In cases of persisting disagreement between the participants of a conversation, the gap view amounts to the no scoreboard view and thus encounters the same problem of violating the intuition of contradiction. Gottschling also contends that the intuition of persisting individual standards is at least weakened on the gap view. Although the truth-values of the scoreboard never contradict the truth-values a knowledge claim would have according to the individual standards, in cases of disagreement the scoreboard's gap value nevertheless does not match the truth-values of the individual standards.

Gottschling finally argues that instead of desperately trying to find a solution that harmonizes incompatible intuitions, we should give up the intuition of contradiction since this intuition is based on a

misunderstanding of contextualism. According to Gottschling, contextualists should be interpreted as claiming that truth-values can vary because the content of the knowledge ascribing sentence can vary in different contexts. So, the skeptic and the commonsensian don't really contradict each other. The only acceptable reading of the intuition of contradiction is, according to Gottschling, that both participants *believe that they have expressed contradictory propositions*. But on this weak reading, the intuition of contradiction is not only fulfilled on the gap view, but also on the different scoreboard and no scoreboard views. Thus, the gap view has no advantage over these other views.

5. ALTERNATIVES TO CONVERSATIONAL CONTEXTUALISM: INFERENTIAL
(OR ISSUE) CONTEXTUALISM, EPISTEMIC CONTEXTUALISM,
AND VIRTUE CONTEXTUALISM

Conversational contextualists accept that skeptical hypotheses are intelligible and that, *if* they are introduced into a given discourse, they are to be taken seriously. Other proponents of contextualist accounts of knowledge try to block skeptical arguments in a more radical fashion. Wittgenstein (1969), in *On Certainty*, has famously challenged the idea that certain ordinary empirical propositions are proper objects for skeptical doubt. His reason, however, is *not* that we know infallibly that they are true. Rather, the idea is that, in ordinary contexts, propositions such as "Here is a hand" are suitable objects neither of knowledge claims nor of the negations of such claims. Instead, Wittgenstein describes them as belonging to the "framework" of our ordinary epistemic discourse, or to our *Weltbild*, without which epistemic discourse would be unintelligible. Wittgenstein concedes that propositions that belong to the extra-epistemic framework of one kind of discourse may in other circumstances be called into question. Calling such propositions into question, however, would mean entering an entirely different kind of language game. The question is whether there is *any* context or language game that can be invaded by global skeptical hypotheses.

Wittgensteinian contextualism, as we may call it, has been quite influential. In particular, it has inspired Michael Williams's highly influential diagnosis and subsequent rejection of skepticism. In his paper "Knowledge, Reflection and Sceptical Hypotheses," Williams calls his position "issue contextualism" (hereafter, IC).¹⁴ IC differs from CC in at least three important respects: First, it is a version of

subject contextualism, according to which the standards a subject *S* must meet in order to know that *p* are set by *S*'s context. Second, in IC the contexts that fix the epistemic standards for knowledge are not primarily conversational; i.e. the mere mentioning of an error-possibility is not sufficient for raising the standards. Rather, the relevance of error-possibilities – and thus the raising of the “level of scrutiny” – depends on *S*'s background information and her practical interests in a specific issue-context. Third, in contrast to CC, issue contextualists deny that there is a single context-independent scale that fixes the degree of epistemic standards in all contexts.¹⁵

For Williams, too, the conversational contextualist's reply to skepticism is too concessive. IC offers a non-concessive anti-skeptical account: In rejecting the idea of a single severity scale for judging epistemic standards, issue contextualists need not hold that skeptics merely raise the standards for knowledge and that therefore the difference between ordinary and skeptical contexts lies only in the different degrees of these standards. According to Williams, the skeptic is no longer studying ordinary knowledge, but knowledge *as such*. But then, compared to a context of studying *specific* knowledge, he enters a “disciplinary meta-context” which involves a number of implicit theoretical presuppositions. Williams calls these presuppositions “methodological necessities” and argues that the skeptic doesn't just raise the standards of knowledge but completely *changes the subject*. And only if we can make sense of the epistemological presuppositions of the skeptical meta-context, which abstract “from everything having to do with human life and human interests,” might skepticism be an interesting epistemological project. Williams denies that such abstractions are intelligible or epistemologically appropriate.

A central feature of Williams's account is his “default and challenge” model of justification. According to this model, the justification of a belief that *p* in a given issue-context requires not only that there be no unjustified challenges to *p*, but also that the belief be reliably formed. In his paper “Inferential Contextualism, Epistemological Realism and Scepticism” Thomas **Grundmann** tries to show that this externalist model of justification undermines Williams's commitment to *epistemological anti-realism*, since the justification of a belief now depends on the reliability of the specific cognitive abilities used in forming the belief. Epistemological realism is, roughly, the view that there are clear-cut generic sources of knowledge which, as Williams puts it, “fix our epistemic situation.”¹⁶ The skeptic exploits this assumption and argues that our ultimate sources of knowledge are defective. Hence if we abandon epistemological

realism, Williams contends, we make an important step toward undermining skepticism. However, Grundmann argues that the reliability of the specific cognitive ability involved in a certain inquiry depends in turn on the reliability of a more general type of cognitive ability. Thus it seems that in contrast to epistemological anti-realism, Williams's account implies natural epistemological kinds.

Further, Grundmann argues that Williams's view about the alleged skeptical implications of epistemological realism is wrong. First, he contends that epistemological realism allows for *different kinds of knowledge about the same domain*, for example, perceptual and *a priori* knowledge of the external world. According to Grundmann, it is thus possible that *a priori* knowledge can help counter the skeptical challenge by justifying the belief that our sense perception is reliable. Second, Grundmann questions Williams's opinion that skeptical hypotheses are *genuine defeaters*, since there seems to be no reason for believing that they are true. Grundmann finally objects that Williams's contextualism cannot adequately deal with the *epistemic meta-inconsistency* in which in context C2 an unrefuted defeater against a methodological necessity of another context C1 arises.

Antonia **Barke** offers an account that is similar in some respects to Williams's IC. She calls her approach *epistemic contextualism*. Like other critics, she argues that conversational features cannot be the sole driving force behind context changes. For Barke, error-possibilities must instead be motivated by our *epistemic inquiries*. A context is defined by the *methods* we use to achieve the goal in question and by the *assumptions* we have to make in order to carry out the inquiry. Context changes are induced by *questioning one or more of the assumptions that underlie the inquiry*. If we cannot defend the questioned assumption by giving independent reasons for its legitimacy, we either have to investigate the assumption in a new inquiry or give it up and change the method so that it no longer depends on this assumption. Either way, the context has changed. According to Barke, her inquiry-driven epistemic contextualism is superior to CC for many reasons. In particular, it can explain the apparent asymmetry between the lowering and the raising of epistemic standards: Once an assumption has been challenged, we have to react to the challenge. Since our epistemic investigations based on this assumption are now called into question, we cannot just forget about the challenge. Only an appropriate answer to it would enable us to return to the previous context – but finding an appropriate answer is usually much harder than attacking an assumption. Another supposed

advantage of epistemic contextualism is its ability to explain a certain feeling of circularity in the zebra case. This feeling arises when we arrive at the belief that the animals are zebras by *assuming* that they are not painted mules, and now conclude from this belief that they are not. Only if we have arrived at the zebra belief through an inquiry that does not depend on the assumption that the animals are not painted mules is it legitimate to deduce “the animals are not painted mules” from “they are zebras.” These considerations also lead, according to Barke, to a clear criterion that differentiates cases in which deduction yields knowledge from those in which it does not.

In his commentary on Barke’s paper, Frank **Hofmann** agrees with Barke’s objections to CC, but doubts whether epistemic contextualism fares any better. His critique consists of three worries. First, since in Barke’s account the epistemic subject has to *believe* that the conditions for deploying a method within an epistemic inquiry are fulfilled, her epistemic contextualism presupposes epistemic internalism and is therefore unacceptable for externalists. A worrisome consequence of Barke’s commitment to internalism, Hofmann argues, is that all epistemically circular arguments turn into *logically* circular arguments. For an externalist, epistemically circular arguments can provide a legitimate source of justification. In externalist accounts *S* can be justified (or warranted) in believing *p* on the basis of a reliable method *M* without *having to believe* that *M* is reliable. It is therefore in turn possible – without getting involved into a vicious logical circularity – to use *p* to justify the reliability of *M*.¹⁷ Since in internalist accounts *S* has to believe that a method is reliable in order to be justified in believing that *p* on the basis of that method, using *p* to justify the reliability of the method would indeed result in an illegitimate form of logical circularity.

Second, Hofmann claims that epistemic contextualism isn’t really an anti-conversational alternative to CC at all. According to Hofmann, in Barke’s account the challenging of an assumption, like the mentioning of an error-possibility in CC, is a *contingent* matter. The dynamics of context changes are thus *not epistemically motivated*. Furthermore, Hofmann contends, epistemic contextualism leads to the same problem of “upward stickiness of epistemic standards” as CC, in that once the standards are raised by a challenge, they tend to stay raised. In particular, since to date the challenge that sense perception is unreliable has not been met, epistemic contextualists need to *ignore* or *forget* this challenge in order to proceed with their epistemic inquiries. For Hofmann, these similarities between epistemic contextualism and CC reveal that *both* accounts regard knowledge as

a *dialectical* phenomenon – and epistemic contextualism is thus nothing but a “revised version of CC.”

Hofmann’s third worry concerns the plausibility of contextualism in general. For Hofmann, contextualism does not account for the intuition that the truth value of “*S* knows that *p*” depends on the fulfilling of certain (objective) conditions which justify the claim that knowledge has been attained. The fulfilling of these conditions, according to Hofmann, is independent of the contingent matter that someone has challenged an assumption or mentioned a skeptical hypothesis. Knowledge is thus a “robust phenomenon” which is context-independent.

In his paper “A Different Sort of Contextualism” John **Greco** proposes an alternative to CC that he dubs “virtue contextualism.” Virtue contextualism maintains that knowledge is true belief resulting from *intellectual* virtue; i.e., true belief that is non-accidental and produced by intellectual excellences such as sound reasoning, correct memory, and accurate perceptions. Greco offers two main reasons for favoring this theory over CC. First, the virtue epistemologist’s account of knowledge can be grounded in a more general theory of virtue and credit, namely, in the theory of moral virtue and moral blame developed by Joel Feinberg (1970). Second, virtue contextualism provides a superior solution to the lottery paradox. Greco maintains that Cohen’s proposed solution is not satisfying, since the salience of the chances of error, i.e. the salience of the possibility of winning the lottery, cannot explain why the overwhelming statistical evidence does not provide sufficient grounds for *S*’s knowing that she will lose. Furthermore, Greco finds it counterintuitive that once the standards have been raised in the lottery case, *S* loses all her knowledge – for example, her knowledge that she bought a lottery ticket or her knowledge of where her car is parked. Virtue contextualism, according to Greco, provides a better explanation of the widely held intuition that *S* does not know that her ticket will lose: Even if the chances of winning the lottery are extremely low, it is nevertheless a matter of *luck* that *S*’s belief that she will lose the lottery turns out to be true. But (salient) luck *undermines credit*. Since *S*’s lucky true belief does not result from her intellectual abilities, she does not know that she will lose the lottery. According to virtue contextualism, this failure of knowledge does not imply that *S* loses all her (fallible) knowledge, since in other cases salient luck need not be involved. Greco therefore contends that virtue contextualism can also explain the apparent difference between the lottery case and the newspaper case, where *S* infers from the results published in a (reliable) newspaper that she has

lost the lottery. In the newspaper case, the truth of her belief that she lost the lottery is not just a matter of chance. It is grounded in and thus credited to her intellectual abilities.

In his paper “On the Prospects for Virtue Contextualism” Dirk **Koppelberg** questions both of Greco’s reasons for adopting virtue contextualism. Koppelberg finds Greco’s account of intellectual credit unsatisfactory for three main reasons. First, the account is formulated only in terms of necessary conditions for deserved intellectual credit. Second, the supposed analogy between Feinberg’s account of moral credit and Greco’s account of intellectual credit is dubious. And third, the distinction between cognitive abilities and intellectual virtues, on which Greco predicates his account of deserved intellectual credit, is blurred.

As far as the lottery paradox is concerned, Koppelberg contends that it is far from obvious that Greco’s solution in terms of virtue contextualism is superior to Cohen’s conversational contextualist solution. In Koppelberg’s estimation, both accounts provide different but equally plausible explanations for the alleged intuition that *S* does not know that she will lose the lottery: Greco explains this intuition by reference to the “statistical nature of the lottery,” and Cohen by the “statistical nature of our reasons.” But, Koppelberg avers, Greco has not given compelling reasons as to why salience of luck should be a better explanation of *S*’s not knowing that her ticket will lose than salience of chances of error. Contra Greco, Koppelberg claims that Cohen’s contextualism can account for the difference between the lottery case and, e.g., the case of the parked car: In high-standards contexts where the chances of error – winning ticket, stolen car – have become salient, *S* fails to know that she will lose the lottery and also fails to know that her car is where she parked it. So, according to Koppelberg, with regard to solving the lottery paradox, virtue contextualism’s prospects are no better than CC’s.

In his paper “Lotteries and Contexts” Peter **Baumann** also discusses Cohen’s contextualist solution to the lottery paradox. Like Greco, he contends that the salience of the chances of error is neither necessary nor sufficient for *S*’s “loss” of her knowledge that she will lose the lottery. According to Baumann, the salience of the (unlikely) possibility that the ticket will win does not account for the widespread intuition that *S* does not know that she will lose the lottery. The explanation is rather that *S* fails to meet the following principle:

- (EP) If it is possible to know that *p*, then there are both good and bad epistemic positions with respect to the proposition that *p*.

Baumann thinks that EP provides the real explanation of the intuition that *S* does not know that she will lose the lottery: *S*'s epistemic position with respect to the lottery proposition (the proposition that she will lose the lottery) is fixed. Since the outcome of the lottery is a pure matter of luck, there is nothing *S* can do to improve (or spoil) her evidence for this proposition. So, according to Baumann, there are no good or bad epistemic positions with regard to this proposition. He also maintains that, in contrast to Cohen's account, his own account can explain the differences between the lottery paradox and the paradox of the preface: Unlike in the lottery case, in the preface case an author can have better or worse evidence with respect to each individual proposition in her book. According to Baumann, fulfilling EP's knowledge condition is required only for knowledge by higher standards. We sometimes use less stringent standards for knowledge according to which it is true to say that *S* *does know* she will not win the lottery. In this weaker sense of "know," "*S* knows that *p*" can be true even though there are no good or bad epistemic positions for *S* with respect to *p*. One consequence of Baumann's account is that *epistemic closure only holds for this weaker sense of knowledge*, where *S* knows both that she will not win the lottery and that she will never be a multi-millionaire. Where higher-standards knowledge is concerned, however, epistemic closure fails, since in such standards *S* does not know that she won't win the lottery, although she does know that she will never be multi-millionaire.

For **Cohen**, one serious weakness of Baumann's account is its (partial) rejection of the principle of epistemic closure. Cohen further argues that Baumann's alternative solution to the lottery paradox is not superior to his own. In particular, Cohen argues that, even if we grant EP's truth, EP still does not explain *S*'s (alleged) failure to know that she will lose the lottery. Moreover, Cohen contends that EP is ambiguous since it remains unclear what it means for an epistemic position to be "good or bad for *S* with respect to *p*." If it means "being in a better or worse epistemic position with respect to *p* than somebody else," then the lottery case fails to meet EP: Knowing about the odds, for example, puts *S* in a *better* epistemic position, compared to somebody who has no idea about the odds of the lottery. But even if EP is interpreted such that there need to be good and bad positions for *S* with respect to knowing that *p*, the lottery case, according to Cohen, fails to meet EP: Compared to the situation in which *S* has not heard about the results, *S* would be in a *better* epistemic position with regard to *p* if she witnessed the drawing of the lottery tickets. Cohen also maintains that his account is quite capable

of explaining the difference between the lottery paradox and the paradox of the preface. In the lottery case the chances of error are salient with respect to the *specific* proposition that *S* will lose the lottery, whereas in the preface case there is only a general worry that one out of many propositions could easily be false.

6. PRIVILEGED CONTEXTS, DEFEASIBILITY, STRENGTH, AND STABILITY

In their paper “Defeasibility and the Normative Grasp of Context,” Mark **Lance** and Margaret **Little** explore the role of (a normative conception of) “privileged contexts” in relation to key epistemological concepts such as knowledge and justification. Their point of departure is a brief consideration of moral particularism, which claims, very roughly, that the moral valence of reasons which count in favor of an action in one context may not count in favor of that same type of action in another context, and that moral thinking should not be construed as an application of general moral principles to particular cases. Now sometimes this position is offered in a version which advocates abandoning all generalizations entirely. This however is misguided, the authors argue. Instead, they say, what is called for are *defeasible generalizations*. The paper defines defeasible generalizations as generalizations that are genuinely explanatory, yet essentially exception-laden. Such generalizations, Lance and Little argue, are crucial for epistemology: Appearances are trustworthy – given that they occur in epistemically privileged conditions. The future will be like the past – except when the conditions for an event’s occurrence deviate in significant ways from the conditions which currently obtain. Their paper’s aim is to explore, on the basis of such considerations, the logical connections between the concepts of privileged epistemic conditions and defeasible generalizations.

Lance and Little are not directly relating their discussion to current forms of contextualism. Yet they consider their position to be a form of epistemic contextualism. It may also be classified as a sort of *epistemic particularism*. Moreover, the approach may have some potential for dealing with skeptical problems. One of its implications, for instance, is that the conditions under which appearances are deceptive can only be explained by appeal to privileged conditions in which appearances are reliable, and that we are entitled to take appearances to be “defeasibly trustworthy.”

Do we really need defeasible generalizations in epistemology? Nikola **Kompa**, in her paper “Moral Particularism and Epistemic

Contextualism,” argues that we don’t. The moral particularist’s view that the reason-giving force of moral considerations is context-dependent involves a *holistic conception of reasons*: In this case, the view is that the question of what in a given situation counts as a good or legitimate reason cannot be answered in terms of general *moral* principles, but only with regard to the entire range of morally relevant features of the context at hand. Lance and Little emphasize strong analogies between moral particularism and their version of epistemological contextualism, which makes them endorse a holistic conception of *epistemic* reasons. However, Kompa argues, the tie between epistemic particularism and epistemic holism can be severed. Moreover, she maintains that, *even if* we embrace a holistic approach to epistemic reasons, this does not necessarily involve commitment to defeasible generalizations. Kompa suggests an alternative form of contextualism that employs, instead of defeasible generalizations, a theory of epistemic *default entitlement*.

In his paper “Stability, Strength and Sensitivity: Converting Belief into Knowledge,” Hans **Rott** focuses on several epistemic concepts that play a fundamental role in certain contextualist approaches to knowledge. In particular, his concern is with explicating the notions of *stability*, *strength*, and *sensitivity to truth*. According to *stability* accounts of knowledge, most prominently championed by Keith Lehrer, *S*’s belief that *p* is an instance of knowledge if and only if *S* would not withdraw this belief on the basis of any true information she might receive. Nozick’s sensitivity-based truth-tracking account of knowledge is tightly connected with the stability account. In particular, the fulfillment of the second subjunctive conditional in Nozick’s definition of knowledge (i.e. if *p* were true, then *S* would have believed that *p*) seems to imply the persistence of *S*’s belief that *p* – even after a conversation with a truthful critic (in a nearby world) who confronts *S* with potential defeaters for *p*. There are also cases, as Rott points out, in which knowledge according to the stability account implies knowledge according to the sensitivity account: The falsehood of Nozick’s first subjunctive conditional (i.e. if *p* had been false, *S* would not have believed that *p*) implies that there is a close not-*p* world in which *S* believes that *p*; i.e. it implies that *S* might believe that *p* when *p* is false. This, in turn, can lead to an instability in *S*’s belief that *p*, since a skeptic could talk *S* out of her belief that *p* by confronting her with the true modal proposition that she *might* be in a position where she is mistakenly believing that *p* is true.

DeRose explicates knowledge in terms of the *strength of the subject’s epistemic position* with respect to the belief in question. If (as

Nozick assumes) the conversation with the critic occurs only in a world that is close to the actual world, the strength of an epistemic position with respect to a true belief p implies its stability. The converse, however, does not hold. Regarding the relation between epistemic sensitivity and strength, DeRose has argued that, as far as believing an ordinary proposition is concerned, a good epistemic position suffices for the belief to constitute knowledge. Regarding belief in the negation of some skeptical hypothesis, however, the epistemic position must be excellent. For the converse, Rott shows that there could be a situation in which a subject S knows that O , but does not know the negation of H according to the sensitivity account, whereas S 's epistemic position with respect to the negation of H is stronger than with respect to O .

In addition to clarifying how the concepts of stability, strength and sensitivity are interrelated, Rott points out a serious problem for the stability account of knowledge: A critic can talk S out of believing a true proposition p and believing a false proposition q by confronting S with the correct information that the conjunction of p and q is false. If S has stronger reasons to believe q than p , she will withdraw her belief that p – and is therefore no longer in a position to know that p . This problem raises some doubts about whether stability accounts provide the right foundation for an adequate explication of knowledge. Finally, Rott calls for a belief revision theory that is also applicable to mainstream epistemological theories based on the notion of *justification*.

In her commentary on Rott, Lydia **Mechtenberg** focuses on two problems facing the stability theory of knowledge. The first problem concerns the counterintuitive result that in cases where all of S 's beliefs are true, none of her beliefs will ever be subject to a belief-revision process as a result of a dialogue with a truthful critic. Mechtenberg contends that this problem is not a genuine problem for the stability account *per se*, but rather only for the standard AGM (Alchourrón, Gärdenfors, Makinson) belief revision theory that serves as a formal model for the stability account in Rott's paper. Mechtenberg objects that this model cannot accommodate all our intuitions regarding stable beliefs, since it only allows *logical* conflicts between S 's beliefs and new information to result in a belief revision. Mechtenberg therefore opts for a revision theory which allows new information about the *improbability* or *implausibility* of S 's beliefs to induce belief contraction.

Mechtenberg's second major worry concerns the "stability problem," namely, that when a subject S has a well-entrenched *false* belief

that p , it is fairly easy for a truthful critic to talk her out of some other true belief that q which is less entrenched than p . Confronted with a true proposition r which contradicts an immediate implication of p and q , S will be willing to give up the true belief that q instead of p . But, as Mechtenberg claims, this stability problem can be mitigated within the framework of a certain *contextualist account of believing* which allows S to assign the conflicting beliefs to *different contexts*. In light of the new information r , it is thus possible for S to postpone her decision between p and q , since she might get further information that will help her make a safe decision according to which no true belief gets lost.

In the remainder of her paper, Mechtenberg outlines her own contextualist account of believing. According to this account, S believes that p *in a given context* if and only if (i) it is to some degree risky for S to act according to her belief that p and (ii) S is ready to incur the risk associated with p in the given context. This context-dependent believing – which Mechtenberg dubs “believing*” – forms the basis for her definition of “believing something in general”: S believes that p if and only if in at least one empirically possible context, S believes* that p . Mechtenberg finally claims that a belief revision theory should incorporate such a contextualist account of believing in order to deal with the stability problem.

ACKNOWLEDGMENTS

Except for the article by Bruce Russell and this introduction all contributions to this issue originate from papers presented at the conference “Contextualist Approaches to Epistemology” held at the Johannes Gutenberg-University of Mainz, Germany, from September 4 to 6, 2003. This conference was supported primarily by a grant from the *Deutsche Forschungsgemeinschaft*. Financial assistance was also provided by the *Gesellschaft für Analytische Philosophie (GAP)* and the *Freunde der Universität Mainz e.V.* We would like to thank each of these organizations for its generous support. Thanks to Kluwer Academic Publishers, the journal *Erkenntnis*, and Hans Rott, its chief editor, for publishing this special issue on contextualism. We wish to thank Verena Gottschling for her help in editing this issue, and Katherine Munn for her help with the language of the present paper. Finally, we would like to thank all of our referees for their valuable help with the project.

NOTES

- ¹ See, for example, Dretske (1970, 1971), Stine (1976), and for a recent defense of an “Expanded Relevant Alternatives Principle,” Heller (1999).
- ² Related views have been articulated by Alvin Goldman (1976) and Peter Unger. Unger says that, if someone utters for example “John knows there is milk on the rug,” the truth value of his utterance depends on the range of contextually *relevant competitors* (Unger, 1984, p. 47f.).
- ³ This label has been proposed by Duncan Pritchard (2002, p. 20).
- ⁴ See Nozick (1981, Chapter 3).
- ⁵ This is the simplest and most straightforward version of the principle of epistemic closure. It has been argued that this version has to be refined into the following version: If *S* knows *p* and knows that *p* entails *q*, and believes *q* as a result of believing *p* and of believing that *p* entails *q*, then *S* knows that *q*. This refined formulation explicitly excludes the case that *S* fails to know that *q* because she does not base her belief that *q* on her knowledge that *p* and that *p* implies *q*. For a discussion of different versions of the principle of epistemic closure see for example Brueckner (1985), Hales (1995) and Barke (2002, Chapter 1).
- ⁶ The most important and influential articles on contextualism by these authors are: DeRose (1995), Cohen (1986, 1988, 1998, 2000), Lewis (1979), (1996).
- ⁷ For reasons of simplicity, in the following we will forebear to mention time *t*.
- ⁸ This problem was first discussed at length by Stephen Schiffer (1996).
- ⁹ Structurally analogous versions of this lottery paradox can be generated with propositions that don’t involve lotteries and don’t appeal to the future. Thus, this paradox is not just a problem for lottery cases or future tense propositions. For example, it seems intuitively correct to say that *S*, having just parked her car in a pretty safe neighborhood, knows where her car is parked. But it seems intuitively wrong to say that *S* knows that her car has not been stolen.
- ¹⁰ Possible solutions to the lottery paradox could consist in rejecting PEC or questioning the intuition that *S* does not know that her ticket will lose. But, as we have already seen, conversational contextualists don’t want to give up the highly plausible principle of epistemic closure. Cohen also takes it for granted that most people have the intuition that *S* does not know that her ticket will lose and that a solution of the lottery paradox along the lines of CC must account for this intuition.
- ¹¹ This criticism has been endorsed for instance by Hilary Kornblith (2000). Kornblith argues, furthermore, that DeRose’s theory only addresses “High Standards Skepticism” regarding knowledge, but fails to address “Full-Blooded (Cartesian) Skepticism,” which maintains that we even have no degree of *justification* whatever for our claims about the external world.
- ¹² Note that what we have called *conversational* or *semantic contextualism*, Engel calls *ascriber-sensitive* or *epistemic contextualism*. This latter term we will reserve for a different position, which will be laid out below. Engel uses “semantic contextualism” to describe the general fact that language is highly context-sensitive.
- ¹³ Stephen Schiffer has also highlighted a number of difficulties with this “error theory” in Schiffer (1996).
- ¹⁴ Elsewhere in the literature, Williams’s account has often been called “inferential contextualism”.

- ¹⁵ The label “issue contextualism” may also be applied to yet another form of epistemological contextualism, i.e. the account proposed in a classic paper by David Annis (Annis 1978). Annis argues for a contextualist theory of *epistemic justification* according to which, relative to one “issue-context,” a person may be justified in believing a given proposition, but not justified in believing this proposition in another context. Annis’s justificatory contextualism is a version of subject contextualism, and there is a certain affinity here to Wittgenstein’s and Williams’s ideas, since Annis maintains that the justificatory status of a person’s belief depends on “certain social practices and norms of justification” (215).
- ¹⁶ For Williams’s account of epistemological realism and anti-realism, see his detailed exposition of his theory in Williams (1996), especially chapter 3, and (2001, pp. 170–172, 191–195).
- ¹⁷ William Alston, for example, has argued at great length about the justificatory power of epistemically circular arguments (see Alston (1993)).

REFERENCES

- Alston, W.: 1993, *The Reliability of Sense Perception*, Cornell University Press, Ithaca.
- Annis, D.: 1978, ‘A Contextualist Theory of Epistemic Justification’, *American Philosophical Quarterly* **15**, 213–219.
- Barke, A.: 2002, *The Closure of Knowledge in Context*, Mentis, Paderborn.
- Brueckner, A.: 1985, ‘Skepticism and Epistemic Closure’, *Philosophical Topics* **13**, 89–117.
- Cohen, S.: 1986, ‘Knowledge and Context’, *The Journal of Philosophy* **83**, 574–583.
- Cohen, S.: 1988, ‘How to Be a Fallibilist’, *Philosophical Perspectives* **2**, 91–123.
- Cohen, S.: 1998, ‘Contextualist Solution to Epistemological Problems: Scepticism, Gettier, and the Lottery’, *Australasian Journal of Philosophy* **76**, 289–306.
- Cohen, S.: 1999, ‘Contextualism, Skepticism, and the Structure of Reasons’, *Philosophical Perspectives* **13**: Epistemology, 57–89.
- Cohen, S.: 2000, ‘Contextualism and Skepticism’, in E. Sosa and H. Villanueva, (eds.), *Philosophical Issues 10: Skepticism*, 94–107.
- DeRose, K.: 1995, ‘Solving the Sceptical Problem’, *The Philosophical Review* **104**, 1–52, quoted from the reprint, in K. DeRose and T. A. Warfield, (eds.), *Skepticism – A Contemporary Reader*, Oxford University Press, Oxford, New York, 183–219.
- DeRose, K.: 1999, ‘Contextualism: An Explanation and Defense’, in J. Greco and E. Sosa, (eds.), *The Blackwell Guide to Epistemology*, Basil Blackwell, Oxford, 187–205.
- DeRose, K.: 2000, ‘Now You Know It, Now You Don’t’, *Proceedings of the Twentieth World Congress of Philosophy*, Vol. V, Epistemology, 91–106.
- DeRose, K.: forthcoming, ‘Single Scoreboard Semantics’, *Philosophical Studies*, prepublication 10/10/2003, 1–21.
- Dretske, F.: 1970, ‘Epistemic Operators’, *The Journal of Philosophy* **67**, 1007–1023.
- Dretske, F.: 1971, ‘Conclusive Reasons’, *Australasian Journal of Philosophy* **49**, 1–22.
- Dretske, F.: 1981, ‘The Pragmatic Dimension of Knowledge’, *Philosophical Studies* **40**, 363–378.

- Dretske, F.: 2004, 'The Case Against Closure', in E. Sosa and M. Steup (eds.), *Contemporary Debates in Epistemology*, Basil Blackwell, Oxford.
- Feinberg, J.: 1970, *Doing and Deserving: Essays in the Theory of Responsibility*, Princeton University Press, Princeton.
- Feldman, R.: 1999, 'Contextualism and Skepticism', *Philosophical Perspectives* **13**, 91–114.
- Goldman, A.: 1976, 'Discrimination and Perceptual Knowledge', *The Journal of Philosophy* **73**, 771–791.
- Hales, S. D.: 1995, 'Epistemic Closure Principles', *Southern Journal of Philosophy* **33**, 185–201.
- Heller, M.: 1999, 'Relevant Alternatives and Closure', *Australasian Journal of Philosophy* **77**, 196–208.
- Kornblith, H.: 2000, 'The Contextualist Evasion of Epistemology', in E. Sosa and H. Villanueva, (eds.), *Philosophical Issues 10: Skepticism*, 24–32.
- Lewis, D.: 1979, 'Scorekeeping in a Language Game', *Journal of Philosophical Logic* **8**, 339–359.
- Lewis, D.: 1996, 'Elusive Knowledge', *Australasian Journal of Philosophy* **74**, 549–567.
- Nozick, R.: 1981, *Philosophical Explanations*, Oxford University Press, Oxford.
- Pritchard, D.: 2002, 'Two Forms of Epistemological Contextualism', *Grazer Philosophische Studien* **64**, 19–55.
- Schiffer, St.: 1996, 'Contextualist Solutions to Scepticism', *Proceedings of the Aristotelian Society* **96**, 317–333.
- Sosa, E.: 2000, 'Skepticism and Contextualism', in E. Sosa and H. Villanueva (eds.), *Philosophical Issues 10: Skepticism*, 1–18.
- Stine, G.: 1976, 'Scepticism, Relevant Alternatives, and Deductive Closure', *Philosophical Studies* **29**, 249–261.
- Unger, P.: 1984, *Philosophical Relativity*, Minneapolis, University of Minnesota Press.
- Williams, M.: 1996, *Unnatural Doubts: Epistemological Realism and the Basis of Skepticism*, Princeton University Press, Princeton, New Jersey.
- Williams, M.: 2001, *Problems of Knowledge: A Critical Introduction to Epistemology*, Oxford University Press, Oxford.
- Wittgenstein, L.: 1969, *On Certainty*, Basil Blackwell, Oxford.

Department of Philosophy
 Johannes Gutenberg University
 55099 Mainz
 Germany
 E-mail: brendel@uni-mainz.de

ELKE BRENDEL

Department of Philosophy
 University of Aberdeen
 Old Brewery
 High Street
 Aberdeen, AB24 3UB, UK
 E-mail: cristoph.jaeger@abdn.ac.uk

CHRISTOPH JÄGER