

Contextualisms in Epistemology

Edited by

Elke Brendel

*Johannes Gutenberg University,
Mainz, Germany*

and

Christoph Jäger

*Aberdeen University, UK, and
University of Leipzig, Germany*

Reprinted from *Erkenntnis* 61 Nos. 2–3 (November 2004)

 Springer

A C.I.P. Catalogue record for this book is available from the Library of Congress.

ISBN 1-4020-3181-5

Published by Springer,
P.O. Box 17, 3300 AA Dordrecht, The Netherlands.

Sold and distributed in North, Central and South America
by Springer,
101 Philip Drive, Norwell, MA 02061, U.S.A

In all other countries, sold and distributed
by Springer,
P.O. Box 17, 3300 AA Dordrecht, The Netherlands.

Printed on acid-free paper

All Rights Reserved

© 2005 Springer

No part of the material protected by this copyright notice may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without written permission from the copyright owner.

Printed in the Netherlands

Table of Contents

ELKE BRENDEL and CHRISTOPH JÄGER / Contextualist Approaches to Epistemology: Problems and Prospects	1
FRED DRETSKE / Externalism and Modest Contextualism	31
CHRISTOPH JÄGER / Skepticism, Information, and Closure: Dretske's Theory of Knowledge	45
MYLAN ENGEL, Jr. / What's Wrong with Contextualism, and a Noncontextualist Resolution of the Skeptical Paradox	61
GILBERT SCHARIFI / Contextualism and the Skeptic: Comments on Engel	91
BRUCE RUSSELL / How to be an Anti-Skeptic and a Non-contextualist	103
WAYNE A. DAVIS / Are Knowledge Claims Indexical?	115
GERHARD ERNST / In Defense of Indexicalism: Comments on Davis	141
VERENA GOTTSCHLING / Keeping the Conversational Score: Constraints for an Optimal Contextualist Answer?	153
MICHAEL WILLIAMS / Knowledge, Reflection and Sceptical Hypotheses	173
THOMAS GRUNDMANN / Inferential Contextualism, Epistemological Realism and Scepticism: Comments on Williams	203
ANTONIA BARKE / Epistemic Contextualism	211
FRANK HOFMANN / Why Epistemic Contextualism Does not Provide an Adequate Account of Knowledge: Comments on Barke	233
JOHN GRECO / A Different Sort of Contextualism	241
DIRK KOPPELBERG / On the Prospects for Virtue Contextualism: Comments on Greco	259
PETER BAUMANN / Lotteries and Contexts	273
STEWART COHEN / Reply to Baumann	287

MARK LANCE and MARGARET LITTLE / Defeasibility and the Normative Grasp of Context	293
NIKOLA KOMPA / Moral Particularism and Epistemic Contextualism: Comments on Lance and Little	315
HANS ROTT / Stability, Strength and Sensitivity: Converting Belief into Knowledge	327
LYDIA MECHTENBERG / The Stability Theory of Knowl- edge and Belief Revision: Comments on Rott	353
List of Contributors	367

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,



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: christoph.jaeger@abdn.ac.uk

CHRISTOPH JÄGER

FRED DRETSKE

EXTERNALISM AND MODEST CONTEXTUALISM

ABSTRACT. Externalism about knowledge commits one to a modest form of contextualism: whether one knows depends (or may depend) on circumstances (context) of which one has no knowledge. Such modest contextualism requires the rejection of the KK Principle (If *S* knows that *P*, then *S* knows that *S* knows that *P*) - something most people would want to reject anyway - but it does not require (though it is compatible with) a rejection of closure. Radical contextualism, on the other hand, goes a step farther and relativizes knowledge not just to the circumstances of the knower, but to the circumstances of the person attributing knowledge. I reject this more radical form of contextualism and suggest that it confuses (or that it can, at least, be avoided by carefully distinguishing) the relativity in what *S* is said to know from the relativity in whether *S* knows what *S* is said to know.

1. INTRODUCTION

I am an externalist about knowledge.¹ Factual knowledge depends on conditions – they are usually external – the knower has (or may have) no knowledge of. I look in the cookie jar and see – and thus know – that there are cookies there. This piece of perceptual knowledge depends on there not being potentially misleading circumstances present. If, unknown to me, people sometimes put fake cookies in this jar, objects I cannot distinguish from real cookies, then, unknown to me, even when there are cookies there, I would not be able to tell, just by looking, that there are cookies in the jar. I would still see the cookies in the jar, but that wouldn't be good enough to know there were cookies in the jar. The knowledge I obtain simply by looking in the jar depends on the absence of a condition I do not know to be absent.

I may, of course believe that there are no potentially misleading circumstances present. And this background belief may be entirely reasonable. Such a belief may be necessary for me to come to believe, just by looking, and, therefore, to see, that there are cookies in the jar. But the point remains: I do not have to *know* that no such condition exists for me to see that there are cookies in the jar. If I had to know that no such condition existed in order to see whether there are



cookies in the jar, there would be precious little, if anything, I could ever come to know by seeing. If there is, unknown to us, a Cartesian demon at work in the universe, deceiving us in random, unpredictable ways, we do not know much, if anything, about the world. Even when we are right (the demon doesn't fool us *all* the time), we don't know. Skepticism is true. Externalism tells us that whether or not skepticism is true, whether or not we know, depends not on our knowing there is no such demon, but on there not being one. That is something we may not know at all. It isn't even clear that we could know it.

Externalism can be expressed in many different ways. The way I have found it convenient to express it is in terms of information.² Knowledge requires information, but whether or not a signal carries information is not itself information the signal carries. You can get the information you need to know without getting the information that what you are getting is information. Instruments that give us information – and, therefore, knowledge – about the quantities they are designed to measure do not provide us with information that what they deliver *is* information. A broken, miscalibrated, or malfunctioning speedometer can look and behave in exactly the way a reliable speedometer behaves. The same is true of the experiences we use in forming perceptual beliefs. We come to know what is happening in our world – we see, hear, and feel what is happening – when our perceptual systems give us information about the world. We don't also have to know that that is what they are giving us. If we had to know this, if, in order to get the information that we were going 60 mph, we always had to get information that the instrument we relied on was, in fact, delivering this information, we would be on the slippery slope to skepticism. We'd never know how fast we were going.

I'm not here going to defend externalism about knowledge. It seems to me pretty obvious. Indeed, if one is going to avoid skepticism, it seems to me inevitable. But, as I say, I'm not here interested in defending this claim. I've done that elsewhere (Dretske, 1969, 1971, 1981). I simply accept this general account of knowledge for purposes of exploring its implications. I'm particularly interested in what it tells us about contextualism.

2. MODEST CONTEXTUALISM

Depending on what one means by *contextualism*, externalism commits one to – indeed, it is simply a form of – contextualism. Whether Clyde, who believes that *P*, knows that *P*, depends on conditions –

the context, if you will – in which Clyde comes to believe *P*. If, unknown to him, the cookie jar he is looking into is a magician's prop, a jar that sometimes (at the whim of his roommate, an amateur magician) contains cookies, at other times wax imitations, Clyde, when peering into the jar, and taking things at face value, can't see what I can see when I look in my cookie jar – that there are cookies there. Evidentially speaking, we are in exactly the same position. We both see cookies in a jar, and we both have exactly the same reasons for thinking that the things we see are cookies. Neither of us suspects, neither of us have reason to suspect, trickery or deception. We both have a true belief – that what we see are cookies. Yet, my belief that they are cookies counts as knowledge, his doesn't. I am in a position to see that they are cookies, he isn't. The difference between us is the context in which we acquire our beliefs, a set of circumstances, favorable in my case, unfavorable in his, that neither of us is aware of. Given his circumstances, there are possibilities Clyde cannot rule out just by looking. This is not so for me. Besides evidence, you need a bit of luck to know.³

I say the difference is context, and it is clear that I mean the context of the person to whom knowledge is attributed – the alleged knower. Whether or not Clyde knows depends on conditions (of which he may be ignorant) existing at the time and place Clyde acquires his belief. If Clyde is using a faulty instrument, he doesn't know even if his belief happens to be true. If he is using a reliable instrument, one that delivers the requisite information, then, other things being equal, he knows. Call this *modest contextualism*. It is so modest, in fact, that some people would probably not call it contextualism. What they mean by contextualism is that knowledge depends not only on the circumstances of the person to whom knowledge is attributed, but on the circumstances of the person who attributes this knowledge. Whether or not *S* knows that *P* is like whether or not *S* is tall: it depends on the context (interests, standards, purposes) of the people who are describing *S*. Are they (Feldman 1999) interested in *S*'s chances of making the basketball team or are they interested in seating arrangements for a dinner party. It makes a difference in whether *S* is tall. According to this form of contextualism, the same is true of knowledge: whether or not *S* knows depends not only on the instruments *S* uses and the circumstances in which *S* uses them, but on what our circumstances (standards, purposes, and commitments) are when we say that *S* knows. If my context is different enough from yours, what I say when I say *S* knows that *P* can be true while what you say when you say it

can be false. I return to this more radical form of contextualism later. For the moment I want only to examine the form of contextualism, *modest* contextualism, to which one is committed by externalism.

Modest contextualism implies the falsity of a doctrine that used to be called the KK Principle. The KK Principle says that if *S* knows that *P*, *S* must know that he knows that *P*. Unlike thirty years ago, there aren't many people around today who accept KK. Too many unacceptable consequences. It denies knowledge to animals and young children simply because they lack an understanding of what knowledge is and, therefore, do not believe (hence, do not know) that they know things. Why should one have to understand what knowledge is in order to see (in the case of children) that Daddy is home or remember (in the case of dogs) where one buried the bone? So no one today, I hope, wants to keep KK. It isn't, therefore, to a theory's great credit that it rejects this doctrine. It is interesting to note, however, that modest contextualism (and, hence, externalism) provides an illuminating explanation of why KK fails. It fails because factual knowledge, according to modest contextualism, depends for its existence on circumstances of which the knower may be entirely ignorant. So the knower can know that *P* without knowing (as required by KK) that he knows that *P*. When he knows that *P*, *S* may, for all he knows, not know that *P*.

Although modest contextualism rejects KK, it does not, taken by itself, deny closure: that one must know all the things one knows to be implied by what one knows. This turns out to be fortunate (for externalism) since although there aren't many advocates of KK around these days, everyone (it seems) wants to defend closure. Closure says that you have to know all the things you know to be necessary for *what* you know (*P*). KK goes beyond this and says you have to know all the things that are necessary for knowing *P* (whether or not you know they are necessary), and, given that knowing *P* implies *P*, that is a much stronger – and therefore less plausible – claim than closure.

I do not, myself, accept closure. I think there are some things we know to be implied by what we know that we do not – perhaps cannot – know to be true. My reason for rejecting closure, however, is not my externalism. It is that I think I have some reasonably clear idea of what kind of evidential relation is required for knowledge (a “conclusive reason” or “information”) and this relation is itself not closed under known implication. We can have, in the relevant sense of “conclusive,” conclusive reasons to believe *P* is true – we can, that is, get information that *P* is true – without having conclusive reasons

to believe, without having information, that Q is true even when we know that P implies Q . I'll come back to closure in a moment when I look at more radical forms of contextualism.

3. RADICAL CONTEXTUALISM

If one thinks that whether or not S knows depends not only on S 's context, but also on the circumstances of the person attributing (or denying) knowledge to S , one is what I will call a *radical* contextualist. Radical contextualism implies that what one says in attributing (denying) knowledge to S can vary from context to context depending on the standards applied (usually implicitly) by the person attributing (denying) knowledge (Cohen, 1986, 1988, 1999; DeRose, 1995; Lewis, 1996; Feldman, 1999).

Radical contextualism provides a quick and snappy answer to skepticism. In a perfectly normal context – in a grocery store, for instance, if someone wonders whether they (pointing) are oranges or tangerines, and I say that I know they are oranges, what I say can be true (so skepticism is false) despite the truth of the skeptic's claim – in the philosophy seminar room – that nobody knows any such thing. We are both right, according to radical contextualism, thanks to different contexts. I attribute knowledge to myself in one context, in the grocery store, the skeptic denies it of me in another, the classroom. Different standards are at work.

What encourages radical contextualism is the plausibility of a relevant alternatives approach to knowledge. If you get it in your head that knowledge that P requires not the ability to rule out or exclude *all* alternatives to P (something we are seldom, if ever, in an evidential position to do), but only the ability to rule out certain relevant alternatives, it appears to be but a short step to relativizing knowledge to the attributional contexts in which different alternatives are deemed relevant. I know they are oranges because, by noting size, shape, and distinctive skin texture, I can visually eliminate the possibility they are tangerines, the only possibility that, given the question ("Are they oranges or tangerines?") seems relevant in the grocery store. When the skeptic ponders the question of whether I know they are oranges, though, that isn't the only relevant possibility. The skeptic wants to know not just whether I can visually distinguish oranges from tangerines (maybe I can), but whether I can distinguish them from wax imitations and hallucinatory figments. If I can't (and I don't know about you, but I can't), then, given the possibilities the

skeptic is worried about (and are therefore salient, if not exactly relevant, in the philosophy seminar room) I do not know they are oranges. At least not in the way I said I knew it in the grocery store – by just looking.

I'm all in favor of focusing on relevant – as opposed to all available – alternatives in assessing whether someone knows something. But the alternatives that are relevant to whether *S* knows that *P*, it seems to me, are always those of *S*'s time, place, and circumstances, not mine. I can't rob *S* of his knowledge by worrying (as he does not) about disembodied brains in vats or by appealing to higher standards. Nor can I create knowledge for *S* by not worrying about, by invoking weaker standards, and therefore not finding relevant, possibilities he does. What the person attributing knowledge to *S* (this can be *S* himself) takes to be relevant is, I submit, irrelevant to whether *S* knows what he is said to know. It is sometimes relevant to what *S* is being said to know, but this, surely, is different from a relevance to whether he knows what he is said to know.

Think about this difference, the difference between what *S* is said to know and whether *S* knows what he is said to know. This difference seems obvious enough, but it is easy to mistake a relativity inherent in the first – what *S* is said to know – for a relativity in the second – whether he knows it. To illustrate with an obvious example think about the use of indexical expressions in describing what someone knows. I say *S* knows that my cousin got married. When I utter those words, I say something true. When you utter those words you say something false. *S* doesn't know that *your* cousin got married. He knows *my* cousin got married. So if we identified what *S* is supposed to know with the words used to express what *S* knows, *S*'s knowledge would be radically contextual. Whether or not he knew would depend on who said it, you or me, and on when we said it.

But this, clearly, is a sensitivity of what *S* is being said to know, not a sensitivity of whether he knows what he is being said to know. One can't argue for a contextualist theory of knowledge on these grounds. One could as well advance a contextualist theory of love and hate because the truth of "Clyde loves (hates) my cousin" depends on attributional context, on who says it and when. Love and hate, as we all know, are relations between persons. What is contextually relative is not Clyde's love or hate for another person, but what person Clyde is being said to love or hate. The same is true of knowledge. What *S* is said to know – the object of knowledge – varies from one context to another with little or no variation in the words we use to describe what *S* knows, but this variation in what *S* is said to know should not

be interpreted as a variation or relativity in whether *S* knows something. The standards for knowledge are – or they may be – constant and invariable from context to context. What varies is what a person is being said to know.

Indexical expressions provide an obvious example, but there are more subtle forms of this contextual variability. I am thinking here (among other things) of contrastive phenomena (Dretske, 1972). I tell you that Clyde knows that Susan stole a scarf. In using these words to describe what Clyde knows, what am I saying that Clyde knows? Am I saying that Clyde knows who stole a scarf – viz., Susan? Or am I saying that Clyde knows what Susan stole – a scarf? Or, perhaps, am I saying that Clyde knows what Susan did with the scarf – she stole it rather than (merely) borrowing it or absent-mindedly taking it home in her purse? Or am I saying – am I *always* saying – all three things? If you and I both use these words to describe what Clyde knows, would we always be saying the same thing? Why couldn't you be describing Clyde as knowing who took the scarf while ignoring (and taking to be irrelevant) questions about whether Clyde knows it was a scarf she took? Why couldn't I, given a different context, be describing Clyde as knowing what she took (he was eyeing that scarf himself) while ignoring questions about whether he knows exactly who took it (Susan or her look-alike sister Sarah). If you and I agree that Susan is the thief, why can't we consult Clyde to find out what she took? Convinced that Clyde knows it was a scarf she stole, we can report what he knows with the words "Susan stole a scarf" even when we know he doesn't know it was Susan. He certainly couldn't pick her out of a police line-up? Why should one have to know it was Susan who stole the scarf to know it was a scarf she stole?

It seems to me fairly clear that in describing Clyde as knowing that Susan stole a scarf we can, depending on where the contrastive focus falls, and the context generally determines where it falls, be reporting at least three different pieces of knowledge. We can be reporting knowledge of the thief's identity, knowledge of the item stolen, or knowledge of the act, what a person did to this item. All this, I assume, is pretty familiar fare. Logicians would prefer to say that the proper, canonical, way to express what Clyde knows is by making scope distinctions explicit. If we want to report Clyde as knowing who stole the scarf without committing ourselves about whether he knows what she stole or that she stole it, we should keep the words "stole" and "scarf" outside the scope of the knowledge attribution. What Clyde really knows is something of or about the person who stole the scarf – that it was Susan. If, on the other hand, we want to describe Clyde as knowing

what Susan stole without committing ourselves on the question of whether he knows who stole it or whether she stole it, the words “Susan” and “stole” are kept outside the scope of the knowledge attribution. He knows of the item that Susan stole that it was a scarf.

I am happy to accept this regimented way of describing differences in what Clyde is being said to know, but it only makes the point I’m after. When we talk, in our ordinary, everyday, way about what Clyde knows, we don’t bother to make explicit these differences in what we say Clyde knows. When I say that Clyde knows that Susan stole the scarf, I’m saying he knows it was Susan who did it. When you, on the other hand, deny this, you may only be denying that he knows it was a scarf she stole. We are both saying something true, but this is not because standards for knowing something shift from one context to another. It is simply a difference in what I am asserting, and you are denying, he knows. The standards for knowing remain the same: in order to know *P* you must, let us say, eliminate all relevant alternatives, but – and this is important – what is relevant does not depend on attributional context. It depends simply on the circumstances in which Clyde saw Susan take the scarf and what, exactly, Clyde is being said to know. Were there other people in the shop who looked (from where Clyde was standing when he saw her take the scarf) enough like Susan to prevent accurate identification? Were there other articles (besides a scarf) that she might have put into her pocket? Could Clyde, from where he was standing, see that it was a scarf? What changes with context is what alternatives we are describing Clyde as being able to rule out. In one case they are alternatives to Susan, the person who stole the scarf, while in another case they are alternatives to the scarf, the item Susan stole. We don’t need a contextual theory of knowledge to account for this difference. All we need is a better understanding of what people are being said to know when they are described as knowing something.⁴

Contextualists might agree with what I’ve said up to this point. They might agree, that is, that what we are saying of someone when we say they know – that those are oranges, for instance – can be different depending on the context in which we say this. A radical contextualist says that the standards – what it takes for a person to know those are oranges, the alternatives he or she must rule out to know this and, therefore, the truth conditions for knowing it – changes from context to context. In the ordinary context she must eliminate only tangerines. In another (skeptical) context, she must also eliminate wax imitations. If she can, visually, distinguish oranges from tangerines but cannot distinguish orange from wax imitations,

then in one, the ordinary, context described above, she can be truly described as seeing (and therefore knowing) that they are oranges, but in the other, skeptical, context, the denial of this claim is true. She cannot see they are oranges. The denial in the skeptical context is consistent with affirmation in the ordinary context. We are simply affirming and denying different things. That, I will be told, is the essential insight of contextualism however we choose to express it – as a change in what a person is being said to know (as I would have it) or (as a radical contextualist would have it) as a change in standards (truth conditions) for their knowing something.

4. CLOSURE

There is, though, a problem with this maneuver as an answer to skepticism. At least there is a problem for anyone who accepts closure. To appreciate the problem, imagine Clyde in the kind of conditions described above. He can visually distinguish oranges from tangerines but he cannot distinguish them from wax imitations. So if he is in a context in which only tangerines are relevant possibilities – e.g., the grocery store – then he can see, hence know, that, they are oranges. He cannot, however, see that they aren't wax. He cannot, and he knows he cannot, eliminate this possibility on visual grounds alone. So if asked whether he knows they are not wax, he would have to say he doesn't know this. He certainly can't see that they are not wax, and he hasn't examined them more closely to know they are not wax in any other way. So now that the possibility of their being wax has been raised, we are in a context in which Clyde doesn't know something (that they are not wax) he knows to be implied by what he formerly said he knew – that they are oranges. So, if closure holds he no longer knows that they are oranges. Asking Clyde whether he can see whether they, the things he knows to be oranges, are not wax imitations, destroys the knowledge he had. He no longer knows what he knew before he was asked this question.

Clyde no longer knows they are oranges because he does not know they are not wax and this is something that, according to closure, he must know to know they are oranges.⁵ Simply by raising this question, we have moved to a context, a skeptical context, in which Clyde not only doesn't any longer know they are oranges, but a context in which he can truly deny having known it before we asked the question. We've moved to a context in which (it is true to say that) Clyde *never* knew they were oranges. The same thing happens if we transport Clyde to a philosophy

classroom, a context in which skeptical possibilities are entertained which Clyde cannot distinguish from the perception of real oranges. We put him into a context in which he can truly affirm not only that he doesn't know there are physical objects and, therefore, doesn't know that there are oranges, but that, in apparent⁶ contradiction to claims he made earlier in the grocery store, he never knew it. According to radical contextualism, thanks to a change in context, his claim today (in a skeptical context) that he didn't really know it yesterday is consistent with his claim yesterday (in an ordinary context) that he knew it.⁷

This is an unfortunate result. It shelters Clyde's knowledge by insulating him from serious skeptical questions. As soon as skeptical questions are raised, Clyde not only ceases to know the things he said (and thought) he knew, he must now agree with the skeptic (in apparent contradiction to what he formerly said) that he never knew these things. If this is an answer to skepticism, it is one that is only available to those who don't think about skepticism. If skepticism is false, let us hope it is false in the philosophy classroom as well as the grocery store. If it isn't, we don't really have an answer to skepticism – at least not one we can give to the skeptic.

What generates this unfortunate result is the adherence to closure. Skeptical questions about whether you can tell oranges from wax imitations destroys your knowledge that they are oranges if, as closure tells us, you have to know (or be in a position to know) that they are not wax imitations in order to know they are oranges. Abandon closure and these unpalatable results vanish. Once I give up closure I can admit (whether I'm on the street or in a philosophy classroom) that I can't (and never could) distinguish, not on straightforward visual grounds, real oranges from perfect wax imitations. But so what? I never said I could. All I said, to someone who wanted to know whether they were oranges or tangerines, was that I could see that they were oranges. What I said I knew – that they are oranges – implies, and I know it implies, that they aren't wax, but if saying, in that ordinary context, that I know they are oranges is consistent with not knowing they aren't wax, then my not knowing they aren't wax, both then or now, is irrelevant to whether I knew they were oranges. Agreeing with the skeptic, in the philosophy seminar room, that I don't now, and never did, know they aren't wax leaves me (unlike a radical contextualist) free to insist that I nonetheless knew what I then said I knew – that they were oranges. I am free, that is, to deny skepticism to the skeptic in a context in which skeptical alternatives are relevant. *That*, it seems to me, is a meaningful answer to skepticism.

But this is an answer to skepticism that depends on the plausibility of saying that what I am claiming I know in the grocery store is something that doesn't require me to know what the skeptic claims I don't (and never did and perhaps cannot) know. It depends for its plausibility on giving up closure. It depends on our willingness to say, for instance, that a person can know that it was Susan who stole the scarf without knowing it was a scarf she stole or even that she stole it. It depends on our willingness to say that a person can see – hence, know – that there are cookies in a jar while not knowing things he knows to be implied by what he knows – that there is a material world, that he is not being cleverly deceived, and that solipsism is false. It depends on our willingness to conceive of knowledge claims in less ambitious terms. As progress reports. Let me take a moment to explain that.

A perceptual report – *S* sees (smells, feels, etc.) that *P* – implies that *S* knows that *P*. It also describes the way *S* comes to know that *P* – by seeing (smelling, feeling, etc.). But, and here is the interesting feature I mean to call your attention to when I urge you to conceive of knowledge in less ambitious terms, these perceptual reports do not tell us how (or even whether) *S* knows many of the things required for *P* to be true.⁸ Perceptual reports are, in this way, progress reports. They describe the way a person got from one (usually unspecified) place to another – the target fact we describe him as knowing. The target fact is *P*, what *S* is said to know by seeing. But the fact or facts from which he came (by visual means in the case of seeing) in reaching *P* is generally left unspecified. And this means you can't evaluate whether *S* could have seen that *P* until you determine where, what fact(s) he came from, in reaching *P*. Suppose *S* is at a wine tasting party. He's sampled each of the wines and is going back for another taste. He sees, and says he can see, that there is still some wine left in bottle #3. Can *S* see what he says he can see? Can he, for instance, see that it not just colored water in bottle #3, something clearly implied by what he said he could see? Clearly not. Should we say, then, that *S* can't really see what he said he could see – that there is still some wine left in bottle #3? Why should we say this? Why not, instead, say that he can see this, but his claim to see it should not be interpreted as a claim to see, or even to be able to see, that it is wine and not colored water. *That* isn't what he was saying he could see. *That* isn't what he was saying he found out by visual means. He was only describing the fact at which he arrived, the target fact, what he came to know by seeing: that there is wine in bottle #3. He wasn't disclosing where he came from in reaching this destination. If he

knows it is wine (and not just colored water) at all, he probably knows this because he tasted it earlier. Or he is simply taking this for granted (it is, after all, a wine tasting party). His current claim – the claim to see that there is wine in the bottle – is knowledge that builds upon that earlier knowledge or set of assumptions. He is claiming to have gone from one place – that the liquid in these bottles is wine – a place he reached by tasting or taking people's word for it – to another place, the target fact – that there is still wine left in bottle #3 – by visual means. He isn't claiming to have arrived, or to even being able to arrive, at the starting point – that it is wine, not colored water – by visual means. That would be a misunderstanding of what he is claiming to have found out by visual means when he said, in that context, that he could see that there was wine in the bottle.

The same is true of Clyde and the cookie jar. Clyde's claim to have seen that there are cookies in the jar implies that Clyde knows there are cookies in the jar, but it does not say how Clyde knows that they aren't fake cookies or merely figments of his own imagination – something implied by what he says he can see. It does not say how Clyde knows that there is a material world, objects that exist independent of his perception of them. Cookies are physical objects, yes, and Clyde knows this, but his knowledge that there are cookies in the jar is not hostage to his knowledge that there is a material world. He certainly can't see that there is a material world. Clyde's claim to have seen that there are cookies in the jar is, normally, a very modest claim. Like our wine example, it is sometimes merely the claim to have found out, by visual means, of the cookie jar, that it isn't (yet) empty (of the things he earlier determined, by tasting, were cookies). You can't tell what Clyde is claiming to have done, or what he is really claiming to know, until you know what he knew before, perhaps by other means, and his perceptual claim doesn't reveal that.

This is why perceptual claims – knowledge claims in general – are best understood as progress reports. They tell the listener where the claimant is supposed to have arrived. They do not reveal where the knower came from in reaching that destination. Knowledge-that-*p* claims are, for this reason, like walked-to-*X* claims. Whether a person can walk to New York City depends on where he walked from in getting there. Was it Boston? Paris? Hackensack, New Jersey? Maybe he couldn't have walked to New York City from Paris. Maybe he could have walked there from Hackensack. The claim that *S* walked to New York City – being, as it were, a progress report – doesn't specify where *S* walked from in getting to New York. It could have been a difficult – maybe even impossible – feat or a leisurely stroll

across a Hudson River bridge. You can't tell. It would be silly, therefore, to object to the possibility of walking to New York City on the grounds that no one can walk on water. Maybe the impossibility of walking on water means one cannot walk to New York City from Paris, but that doesn't mean a person (even a Parisian) cannot walk to New York City. They can walk there from Hackensack. No trick at all.

Exactly the same is true of knowledge. Whether Clyde can see that there are cookies in the jar depends on where he came from in reaching that conclusion. If all he is claiming to have done (and there are many conversational contexts in which this is all one would be claiming to have done) is to tell, visually, that the jar isn't empty of things that are known (or assumed), perhaps on other grounds (or no grounds at all) to be cookies, then an inability to distinguish real cookies from fake cookies is irrelevant to what Clyde is being said to know.

If, then, we take knowledge claims as progress reports, reports of whether – and, if so, how – we arrived at the target fact from places unspecified, there is the same kind of contextual variability in what we are saying we know as there is in what we are saying we did when we say we walked to New York City. It is as silly to object, on skeptical grounds, that one can't know there are cookies in the jar or oranges on a shelf as it would be to insist, for instance, that Parisians cannot walk to New York City. It exhibits a misunderstanding of what is being said.

NOTES

¹ But not about justification.

² Since Dretske (1981). Before that I expressed it in terms of conclusive reasons (Dretske, 1969, 1971).

³ It can be lucky that *S* knows what they are going to do (he happened, by chance, to be standing nearby when they discussed their plans) without its being at all lucky – a matter of chance – that he is right about what they are going to do. See Unger (1968).

⁴ After writing an early draft of this article I read Heller (1999) who, in slightly different language, says much the same thing. What I don't know about Heller is whether he is willing to embrace the consequences of this view that I discuss below – viz., the denial of closure.

⁵ I assume throughout this discussion that all the relevant subjects understand and know the relevant implications – that, for example, if something is an orange it is not wax.

⁶ It is only apparent. See next sentence.

⁷ Though the truth of what he said yesterday isn't consistent with his saying today that what he said yesterday was false.

⁸ I developed this idea in Dretske (1969).

REFERENCES

- Cohen, S.: 1986, 'Knowledge and Context', *The Journal of Philosophy* **83**, 574–583.
- Cohen, S.: 1988, 'How to Be a Fallibilist', in J. Tomberlin (ed.), *Philosophical Perspectives*, Vol. 2, Ridgeview Publishing, Atascadero, CA, 91–123.
- Cohen, S.: 1991, 'Scepticism, Relevance, and Relativity', in B. P. McLaughlin (ed.), *Dretske and His Critics*, Blackwell, Cambridge, MA, 17–37.
- Cohen, S.: 1999, 'Contextualism, Skepticism, and the Structure of Reasons', in J. Tomberlin (ed.), *Philosophical Perspectives*, Vol. 13, *Epistemology*, Ridgeview Publishing, Atascadero, CA, 57–90.
- DeRose, K.: 1995, 'Solving the Skeptical Problem', *Philosophical Review* **104** (1), 1–52.
- Dretske, F.: 1969, *Seeing and Knowing*, University of Chicago Press, Chicago.
- Dretske, F.: 1970, 'Epistemic Operators', *Journal of Philosophy* **LXVII** (24), 1007–1023.
- Dretske, F.: 1971, 'Conclusive Reasons', *Australasian Journal of Philosophy* **49**, 1–22.
- Dretske, F.: 1972, 'Contrastive Statements', *Philosophical Review* **LXXXI** (4), 411–437.
- Dretske, F.: 1981, *Knowledge and the Flow of Information*, MIT Press, A Bradford Book, Cambridge, MA.
- Dretske, F.: forthcoming, 'The Case Against Closure', in M. Steup and E. Sosa (eds.), *Debates in Contemporary Epistemology*, Blackwell, Cambridge, MA.
- Feldman, R.: 1999, 'Contextualism and Skepticism', in J. Tomberlin (ed.), *Philosophical Perspectives*, Vol. 13, *Epistemology*, Ridgeview Publishing, Atascadero, CA, 91–114.
- Heller, M.: 1999, 'The Proper Role for Contextualism in an Anti-luck Epistemology', in J. Tomberlin (ed.), *Philosophical Perspectives*, Vol. 13, *Epistemology*, Ridgeview Publishing, Atascadero, CA, 115–130.
- Lewis, D.: 1996, 'Elusive Knowledge', *Australian Journal of Philosophy* **74** (4), 549–567.
- Unger, P.: 1968, 'An Analysis of Factual Knowledge', *The Journal of Philosophy* **65**, 157–170.

Department of Philosophy
 Duke University
 201 West Duke Building
 Durham, NC 27708-0743
 USA
 E-mail: dretske@acpub.duke.edu

CHRISTOPH JÄGER

SKEPTICISM, INFORMATION, AND CLOSURE: DRETSKE'S THEORY OF KNOWLEDGE

ABSTRACT. According to Fred Dretske's externalist theory of knowledge a subject knows that p if and only if she believes that p and this belief is caused or causally sustained by the information that p . Another famous feature of Dretske's epistemology is his denial that knowledge is closed under known logical entailment. I argue that, given Dretske's construal of information, he is in fact committed to the view that both information and knowledge *are* closed under known entailment. This has far-reaching consequences. For if it is true that, as Dretske also believes, accepting closure leads to skepticism, he must either embrace skepticism or abandon his information theory of knowledge. The latter alternative would seem to be preferable. But taking this route would deprive one of the most powerfully developed externalist epistemologies of its foundation.

1. EXTERNALISM, INFORMATION, AND THE KK THESIS

Externalist theories of knowledge typically claim to provide an effective response to skepticism. One of the most influential externalist epistemologies has been developed by Fred Dretske. The core idea of his theory is that:

- (K) An epistemic subject knows of some object (or source of information) s that it has the property F if and only if the subject believes s to be F and this belief is caused (or causally sustained) by the information that s is F .

Dretske has spelled out this view in great detail. According to the account of information laid out in his book *Knowledge and the Flow of Information*, a signal carries, relative to a given subject, the information that s is F if and only if the conditional probability of s 's being F , given the signal and the subject's background knowledge, is 1, but, given only the subject's background knowledge, is less than 1.¹

The central feature of this account, which is primarily designed for perceptual knowledge, is that the information relation is veridical:



The appearance of the (shelves in the) refrigerator can only inform you that the refrigerator is empty if it is true that it is empty. The gas gauge cannot carry the information that the tank is half full unless it is half full. Moreover, due to the theory's reference to background knowledge, it may happen that a signal carries some piece of information for me but not for you: if you saw Jenny walking up the stairs to the front door, but I did not, her signature pattern of knocks at the door may inform me that she has arrived, but it does not carry that information for you, since you already know that she is there.

How does skepticism fare in this theory? As Dretske observes, it seems undeniable that skeptical worries undermine our knowledge *that we know* (Dretske, 1981a, p. 128). A crucial feature of his externalism, however, is that – contrary to what Roderick Chisholm, Carl Ginet, and others have argued – knowing that *p* does not require that the subject possess the higher-order knowledge that she knows that *p*. The KK thesis, the claim that in order for *K* to know that *p*, *K* must also know that she knows this, is rejected in Dretske's account. Put in terms of information: if *K* comes to know something by receiving some piece of perceptual information, she need not receive the higher order information *that it is (genuine) information* she receives. I can see that there is beer in the fridge, even if I don't possess any higher-order information about whether – to use Chisholm's idiom – my currently being appeared to beerly is reliable.

The fact that the KK thesis immediately leads into skeptical problems is a good reason for abandoning it.² Yet, can Dretske's account also handle skeptical worries from other directions? Dretske himself was the first philosopher to give clear articulation to the fact that another road to skepticism is the view that knowledge is closed under known logical entailment (see especially Dretske, 1970, 1981b). In a first approximation, the relevant principle of epistemic closure can be sketched as follows:

(PEC) If *K* knows that *p* and knows that *p* entails *q*,
 then *K* knows that *q*.³

Now, Dretske famously rejects closure as well. His arguments for this conclusion are controversial, but this is not what I want to discuss in this paper.⁴ My question is whether this part of Dretske's epistemology is compatible with his information-theoretic account of knowledge. Is his denial of closure even *consistent* with his idea that generating perceptual knowledge consists essentially in receiving information, given that the information relation is construed in the

probabilistic way sketched above? I argue that it is not. There are deep perplexities regarding the compatibility of these two cornerstones of Dretske's epistemology, i.e. his information-based, probabilistic account of knowledge, and his denial of closure.

2. CONTEXTUALISM AND CLOSURE

Dretske's *Knowledge and the Flow of Information* is a modern classic in externalist epistemology, and his seminal work on closure dates back to the 1970s (see, for example, Dretske, 1970–1972). In his paper “Externalism and Modest Contextualism” he takes some fresh looks at skepticism, closure, and knowledge by comparing his externalism with epistemic contextualism. I shall take my cue from these recent reflections of Dretske's.

Contextualist theories of knowledge, as championed in recent decades by Stewart Cohen, Keith DeRose, and David Lewis,⁵ differ considerably in their details. Yet they share the idea that the truth values of knowledge attributions can vary with the epistemic standards of the attributer's context. The hallmark of this “attributorial” or, as Dretske calls it, “radical contextualism” is its contention that it can properly deal with skeptical puzzles, while at the same time retaining the view that knowledge is closed under known logical entailment. Let K stand for some normal human adult, e for some suitable empirical proposition (such as “This is a zebra,” “Here is a hand,” etc.), and not- h for the negation of some (local or global) skeptical alternative to e . The skeptical puzzle arises from the fact that we would neither want to deny that, in a suitable situation, K knows e , nor that K knows that e entails not- h . If we accept that knowledge is closed under known entailment, it follows that K also knows not- h . But skeptical arguments seem to show that the latter is impossible: we don't *know*, it seems, that skeptical hypotheses are false. The problem is that each of these claims, including PEC (or some embellished version of this principle), appears to be true; but taken together, they are inconsistent. Which one has to give?

Very roughly, the contextualist replies that in contexts in which skeptical hypotheses are salient, the standards for knowledge ascriptions differ from the standards for knowledge ascriptions in everyday contexts. For example, in ordinary contexts the reasons we have for taking appearances to be trustworthy suffice to license the judgement that our perceptual beliefs constitute knowledge. In

skeptical contexts, however, standards have been raised, and relative to these elevated standards knowledge attributions become false. Closure is said to hold in both types of context. Hence, says the contextualist, in ordinary contexts we do know after all that skeptical hypotheses are false, whereas in skeptical contexts we don't know that ordinary empirical propositions are true.⁶

These claims are Dretske's main target. According to conversational contextualism, we can apparently rob someone of his knowledge simply by raising the question whether he is in an epistemic position in which he can exclude skeptical alternatives. Moreover it is possible in this view for a subject, just by contemplating some skeptical possibility and thereby making it a citizen of the realm of relevant alternatives to a given proposition, to create a context in which it is true to say that she *never knew* that proposition. This, Dretske urges, is a most unfortunate result. For if this is an answer to skepticism, it is one that is only available to those who never think about skepticism. Yet, if skepticism is false, shouldn't it be false in the philosophy classroom as well as the grocery store?

I think that here Dretske does put his finger on an important problem of contextualism. I want to ask, however, whether his information theory of knowledge is really any better off. Dretske recommends rejecting closure in order to avoid these unwelcome results. "Abandon closure", he writes (2004, p. 182), "and these unpalatable results vanish." This is right. Rejecting closure is a theoretical option that, if adopted, would solve the kind of skeptical problem outlined above without forcing us to embrace a contextualist position. But giving up closure is a very high price to pay. My question is whether it is a price that Dretske himself is able to pay without having to sell his probabilistic theory of information. I shall now argue that it is not.

Dretske offers two main lines of reasoning against closure. One is a family of indirect arguments to the effect that abandoning closure is unavoidable if we want to put skeptics in their place.⁷ It is doubtful, however, whether these arguments can ultimately cut any ice. To begin with, it is immensely plausible to assume that a body of knowledge can expand by deductive reasoning from known premises. Giving up closure would mean giving up this integral part of our ordinary understanding of knowledge. Secondly, Dretske's *modus ponens* is the skeptic's *modus tollens*: the skeptic will adhere to closure, insist on his claim that we don't know that skeptical hypotheses are false, and thus conclude that we lack knowledge of those homely truths we ordinarily take ourselves to know.⁸

Yet in “Externalism and Modest Contextualism” Dretske also suggests another interesting argument. “I do not, myself, accept closure”, he writes:

I think there are some things we know to be implied by what we know that we do not – perhaps cannot – know to be true. My reason for rejecting closure, however, is not my externalism. It is that I think I have some reasonably clear idea of what kind of evidential relation is required for knowledge (a “conclusive reason” or “information”) and this relation is itself not closed under known implication. We can have, in the relevant sense of “conclusive,” conclusive reasons to believe p is true – we can, that is, get information that p is true – without having conclusive reasons to believe, without having information, that q is true even when we know that p implies q . (Dretske, 2004, p. 176f)

The argument in this passage seems to be this: (perceptual) knowledge requires receiving appropriate information from some signal. But the relation of receiving such information is not closed under known implication; hence (perceptual) knowledge is not closed under known implication. In a forthcoming paper on closure that Dretske has kindly given me the chance to read, he explicitly argues that no signal can carry the information that skeptical hypotheses are false. “There is nothing in the world,” he says, “[...] that indicates that there is a material world” (Dretske, forthcoming). Similar remarks can already be found in *Knowledge and the Flow of Information*:

No signal can rule out *all* possibilities if possibilities are identified with what is consistently imaginable. No signal, for instance, can eliminate the possibility that it was generated, not by the normal means, but by some freak cosmic accident, by a deceptive demon, or by supernatural intervention. (Dretske, 1981a, p. 130)

I agree. I wish however to argue, first, that Dretske’s account is – at least with respect to empirical propositions and their antiskeptical consequences – committed to the view that the relation of a signal’s carrying information *is* closed under known entailment. Second, I will show that, given this fact and given that we should work with a refined version of the closure principle, Dretske’s theory is also committed to the closure of knowledge. Contrary to what he believes, his information-theoretic externalism implies, at least for ordinary empirical propositions and their anti-skeptical consequences, that knowledge *is* closed under known entailment.

3. INFORMATION AND CLOSURE

Let us first examine whether *information* is closed under known entailment, i.e. whether it holds that:

- (PIC) If r carries, relative to the subject K , the information that p , and K knows that p entails q , then r carries, relative to K , the information that q .

Recall, first, that in Dretske's account a signal r carrying the information that p (relative to a given subject) implies that the probability of p , given r (and the subject's background knowledge k) is 1. And if K knows that that p entails q , p does entail q . Hence if in addition the conditional probability of p , given some signal r (and k), is 1, the conditional probability of q , given r (and given k), *must* also be 1.

This does not yet amount to an effective counterargument. For as we have seen, Dretske's definition of a signal's carrying information is slightly more complicated: it includes the condition that the prior probability of the proposition in question is less than one. In Dretske's theory, it holds that:

- (I) A signal r carries the information that s is F (relative to a given subject) if and only if $P(F(s)/r \ \& \ k) = 1 \ \& \ P(F(s)/k) < 1$.

Nevertheless, I wish to argue that information in the full sense of this account, at least for the kinds of propositions at issue, is closed, not only under known entailment, but also under entailment *simpliciter*. The propositions at issue are (potential) contents of perceptual beliefs, and their anti-skeptical consequences. Let e stand again for some empirical proposition of an appropriate kind ("This is a zebra," "Here is a hand"), and let not- h represent an appropriate proposition which negates some (local or global) skeptical hypothesis that is incompatible with e ("This is not a cleverly painted mule," "There is a material world"). What needs to be shown, then, is that the following instantiation of the above closure principle regarding information (PIC) is *true*:

- (PIC*) If $P(e/r \ \& \ k) = 1 \ \& \ P(e/k) < 1$, and Knows
 $(K, e \Rightarrow \text{not-}h)$, then $P(\text{not-}h/r \ \& \ k) = 1 \ \& \ P(\text{not-}h/k) < 1$.

In prose: if r carries, relative to K 's background knowledge k , the information e , and K knows that e entails not- h , then r *does* also, relative to k , carry the information that not- h .⁹ I have already indicated the first step of the argument for this claim: if we assume that a signal r carries the information e (which implies that the conditional probability of e , given the signal and k , is 1), and that K knows that e entails not- h (and hence that e does entail not- h), we must conclude

that the conditional probability of not- h , given r and k , is also 1. Put in the lingo of the probability calculus:

- (1) $P(e/r \ \& \ k) = 1 \ \& \ P(e/k) < 1$,
- (2) $e \Rightarrow \text{not-}h$.

Therefore:

- (3) $P(\text{not-}h/r \ \& \ k) = 1$.

But can it also be shown that the antecedent probability of not- h is less than 1, i.e. that $P(\text{not-}h/k) < 1$?

At this point, we might ask whether this second condition is acceptable. Does not this requirement take us too far away from our intuitive understanding of the concept of information? For this requirement has the consequence that, if K already knows that p , it is not possible that there still be any signal for K that carries the information that p . You see Jenny approaching, and a few moments later you also hear her voice. Should we not say that this auditory signal does carry the information for you that she is here, despite the fact that you already know this? Perhaps. At least one may be inclined to say so if the counterfactual holds that, had you not already seen (and thus known) that she is there, hearing her speak to you would have been sufficient for you to generate that knowledge. However, dropping Dretske's condition about the antecedent probability of informational content would amount to a substantive revision of his account. Fortunately such a move is not necessary to bring our point home. For with respect to negations of *skeptical hypotheses*, the requirement that their antecedent probability be less than 1 must, in the framework of Dretske's epistemology, clearly be regarded as fulfilled. The reason is that Dretske takes skepticism seriously. "Skepticism," we even hear him saying, "is true" (Dretske, 2004, p. 174). Now by this, I take it, he does not mean to say that skeptical hypotheses are actually true, but rather that the skeptic is right in claiming that we are not entitled to be certain that they are false. The assumption, in other words, is that (given what we know about the world) the probability of skeptical hypotheses being true is not zero. This means of course that the probability of their negations – i.e. of assumptions of *not* being brains in vats, about the existence of a material world, etc. – is less than 1. So from this part of Dretske's epistemology we get:

- (4) $P(\text{not-}h/k) < 1$.

And hence we finally arrive at:

- (5) $P(\text{not-}h/r \ \& \ k) = 1 \ \& \ P(\text{not-}h/k) < 1$.

(5) says that the signal r carries the information that not- h , and thus that our skeptical hypothesis h , which is incompatible with the empirical proposition e , is false. In summary, then, it has emerged so far that, at least for the kind of propositions here at issue, Dretske's information relation *is* closed under (known) logical entailment: if there is a perceptual signal that carries the information for you that you have hands, and if you know that if you have hands, there is a material world, then that signal also carries the information for you that there is a material world.

What should be concluded from this? Dretske must either take on board the idea that information is closed under (known) entailment, or give up his theory of information. However, each alternative has shattering consequences.

Consider first what would happen if we adopted a Dretske-style theory of knowledge and information, thereby accepting that information is closed under (known) entailment. Someone may want to defend Dretske's overall attack on skepticism along the following lines: "All right," it may be responded, "you have shown that, for perceptual propositions and their anti-skeptical consequences, the *information relation* is closed under known entailment. But what we are really interested in with regard to skeptical puzzles is whether *knowledge* is closed under known entailment. So even if you are right about information, you are barking up the wrong tree. Your argument has little impact on what ultimately is at issue." But this objection would be misguided for several reasons.

First, Dretske seems to think that (1) if information is not closed under known entailment, then neither is knowledge. He further seems to think that (2) the antecedent of this condition is fulfilled (and that therefore knowledge is not closed). But if what has been said so far is right, this argument does not go through since it relies at least on one false premise, namely (2). It is of course another question whether the first premise of this argument is true. If this is false, the abstract possibility remains for Dretske to retain the view that closure fails for knowledge, although (as I have shown) it does not fail for information. It may be noted in this context that Dretske says in the above quotation that we can "*get* information that p is true without *having* information that q is true even if we know that p implies q ." This might be interpreted as referring to a relation other than a signal's merely carrying information. However, could it be a reasonable position within Dretske's epistemological framework to accept that the relation of "carrying information" is closed under known entailment, while denying this for knowledge? I will return to this question

shortly, but first I would like to stress that full-blown knowledge skepticism is in any case not the only problem that we must come to grips with in our epistemological endeavors.

Our discussion thus far already shows that the theory under consideration leads to skepticism about *perceptual information*. And that is bad enough. Consider for instance global skeptical hypotheses. How *could* a perceptual signal relay the information that such hypotheses are false? How could a view of Jenny, nice as it may be, carry the information that there is a material world? If there is one lesson to be drawn from skeptical arguments, it is that perceptual experience certainly cannot teach us that skepticism is false. Put in terms of information: no perceptible signal can carry information about the falsity of skeptical hypotheses. We thus have arrived, to begin with, at a *reductio* of Dretske's theory of information. If we concede, as we should, that perceptual experience is neutral with regard to skeptical scenarios, any account that yields a result to the contrary must have gone astray. I have already indicated that Dretske himself seems to share that view. He speaks of anti-skeptical implications of empirical propositions as "heavyweight implications," and he underlines that:

Ordinary things we come to know by perception always have heavyweight implications that are out of range: we cannot see (hear, smell, or feel) that they are true. I can see that there are cookies in the jar, but I cannot see that there is an external world. [...] This is true of all indicators, all sources of information. That is why there is nothing in the world – either mental or material – that indicates that there is a material world. Nothing in the present that indicates there is a past. (Dretske, forthcoming)

Exactly. But if what I have said is on target, Dretske's information theory of knowledge commits him to precisely these consequences which he himself declares untenable. The problem is that we must either enter the den of skepticism with regard to perceptual information, or develop a theory of perceptual information that differs considerably from the one Dretske proposes. This is an important epistemological result on its own.

4. KNOWLEDGE CLOSURE

Let us now return to knowledge closure. A person knows that *p*, we are told, if and only if her belief that *p* is caused (or causally sustained) by the information that *p*. How can an abstract entity like information act in the world of causes? The metaphysics of causation is an area of persisting controversies. Yet, according to the standard

view, as advocated for instance by Donald Davidson, David Lewis, Jaegwon Kim, and also Dretske himself,¹⁰ causes must be events. Dretske adds that they are events which are effective in virtue of having certain properties (Dretske, 1981a, p. 88). Regarding the relation between perceptual signals and beliefs, we may say that perceptual signals, construed as events, cause beliefs in virtue of the fact that they carry certain pieces of information. Thus, Dretske's account of knowledge may also be represented as follows:

- (K*) An epistemic subject *K* knows of some (perceptual) object *s* that it has the property *F* if and only if there is a signal *r* which carries the information that *s* is *F*, and *r* causes (or causally sustains), in virtue of carrying that information and via a non-deviant causal chain, *K*'s belief that *s* is *F*.

The qualification "via a non-deviant causal chain" is needed to exclude cases in which a signal causes *K* to hold a true belief but produces that belief in the wrong way, as for instance when some neurological instrument, triggered by some signal, produces the belief in *K* by direct brain stimulation.¹¹ Causal theories of knowledge and belief are faced with the task of providing an account of non-deviant causal chains that lead from (potential) sources of epistemic attitudes to the appropriate beliefs. This has proven to be a difficult task. But let us assume, for the sake of the argument, that a satisfying account can be worked out.

Next it should be noted that there are important independent reasons for adding a condition to the simple formulation of the closure principle PEC that we (and Dretske) have been working with so far ("If *K* knows that *p* and knows that *p* entails *q*, then *K* knows that *q*").¹² It is widely acknowledged that, as it stands, PEC is problematic simply because it is doubtful whether *belief* is closed under known (or believed) entailment. On a common-sense understanding of the notion of belief, for instance, it is not thus closed: it is certainly not a conceptual truth that people believe every logical consequence of what they believe, even if they see that they follow. However, what we are after when investigating *epistemic closure* (as opposed to what may be called "doxastic closure") is an illuminating analysis of conditions specific to the concept of knowledge. The principle to be examined should therefore not be vulnerable on account of incorporating a dubious assumption about belief.

The obvious way to protect the principle from this lapse is to add a condition to PEC which says that the subject at least also *believes* q .¹³ However, this manoeuvre does not yet exclude that the belief that q is held for some reason that has nothing to do with the subject's knowledge about p and p 's consequences. Hence we should strengthen the antecedent of the relevant closure principle to the effect that the belief that q be a *result of*, or be *based on*, the subject's knowledge (and beliefs) that p and that p entails q . In summary, then, the epistemologically interesting question is not whether PEC, but whether a principle of epistemic closure along the following lines is true:

- (PEC*) If K knows that p , and knows that p entails q , and believes q on the basis of knowing (and hence believing) p and knowing (and hence believing) that p entails q , then K knows that q .

I believe this principle *is* true, but I shall not argue for this claim here. Stephen Hales (1995) has argued that it is even *trivially* true. Advocates of such a principle, he claims, face the task of showing that it is nevertheless not philosophically empty. The argument I will present now meets this requirement. My argument shows that if the principle is true, certain externalist theories of knowledge fall prey to skepticism. This, I take it, is not a philosophically trivial result. Moreover, the important point in the present context is that, whatever we may think about PEC*, Dretske's epistemology is committed to this principle. And if this is true, and if we accept, with Dretske, that the skeptic is right in claiming that we don't know that skeptical hypotheses are false, Dretske's theory also falls victim to knowledge skepticism.

So let us feed into PEC* the relevant parameters. First, it seems clear that in Dretske's account the notion of believing q "as a result" or "on the basis" of believing p must be interpreted in causal terms. His theory of knowledge clearly suggests a causal interpretation of the epistemic basing relation. Let us tacitly understand that non-deviant causal-chain conditions are fulfilled. Let e (as an instantiation of p) again be some suitable empirical proposition and not- h (as an instantiation of q) one of e 's anti-skeptical consequences. Given a causal interpretation of the basing relation, the antecedent of PEC* then gives us that K 's belief that e , which is part of what constitutes K 's knowledge that e , is one of the causes of K 's belief that not- h . But if K 's belief that e is an instance of perceptual belief, it must, according to Dretske, have been caused by some perceptual signal r which carries the information e . It thus follows that r is also among the causes of K 's

belief that not-*h*. Furthermore, as I have shown in Section 3, if *r* carries the information that *e*, and *K* knows that *e* entails not-*h*, then *r* must also carry the information that not-*h*. Hence *K*'s belief that not-*h* has been caused by a signal that carries the information not-*h* – and isn't that precisely what suffices, according to Dretske's definition, for that belief to be an instance of knowledge?

Prima facie, there is still a fly in the ointment. It may be objected that, strictly speaking, *r* can only be regarded as a partial cause of *K*'s belief that not-*h*. For PEC* rules (in its causal interpretation) that this belief must be caused, via a non-deviant causal chain, by the belief that *e* and the belief that *e* entails not-*h*. Would Dretske nevertheless say that *K*'s belief that not-*h* amounts to knowledge? This is the story he has to tell. I think he would in fact admit that we are dealing with a case of knowledge here. More importantly, he certainly *should* say so. For what could conceivably prevent such a belief from deserving the title of knowledge? By hypothesis, the belief that not-*h* is caused (a) by a signal that carries the information that not-*h*, and (b) by the subject's true belief that her belief that not-*h* is entailed by a true belief she holds (i.e. her belief that *e*). What could be safer than that? What else could be demanded in an externalist theory such as Dretske's?

5. CONCLUSION

I take these reflections to show that it will be preferable for Dretske to opt for the second alternative mentioned at the end of Section 3 and give up his probabilistic theory of information. At least he needs to revise that theory substantially if it is to handle the problems discussed in this paper. Until an alternative is on the table, however, his theory, which is one of the most detailed and most powerfully developed externalist accounts of knowledge, has lost its foundation. Moreover, it cannot be regarded as a viable alternative to contextualism. Dretske is one of the great pioneers of externalism in epistemology, and much of the force and attraction of his externalism has been due to its carefully crafted account of information. We now realize that, as it stands, this theory does not escape skepticism.

ACKNOWLEDGMENTS

For helpful comments and discussions I would like to thank Elke Brendel, Mark Siebel, an anonymous referee for *Erkenntnis*, and

especially Fred Dretske. I cannot resist the temptation to say that Dretske has agreed, in correspondence and in his reply to my presentation of this material at the conference in Mainz, that the arguments developed in this paper are on target. He seems to be fairly optimistic, however, that his theory can be repaired without giving up too much of its central idea. I do not think that the prospects for this are too rosy; indeed I believe that the issues I have discussed here point to a more general and fundamental problem of externalist accounts of knowledge: they avoid higher-order skepticism by rejecting the KK thesis, but they are committed to closure (and hence invite skeptical arguments that work with closure). However that may be, it is always much easier to criticize a theory than to come up with a relevant alternative. I wish I had one.

NOTES

- ¹ Dretske (1981a, p. 65). See also Dretske et al. (1983, p. 57). For his definition of knowledge see (1981a, p. 86), and (1983, p. 58). Instructive reconstructions and discussions of Dretske's information theory of knowledge can be found in Foley (1987), the "Open Peer Commentary" in Dretske (1983), and in the papers collected in McLaughlin (1991).
- ² For various epistemological problems the KK thesis creates see for example Greco (2000, pp. 181–184).
- ³ As I shall discuss below, this simple formulation of the closure principle is vulnerable to objections that have nothing to do with skepticism. I will embellish this formulation in Section 4, but for a start it will suffice to work with PEC. Dretske's argument that closure invites skepticism will be sketched in the next section.
- ⁴ That Dretske's famous Zebra Case does not in fact constitute a genuine counterexample to the closure principle has for example been argued by Jonathan Vogel (1990). Yet, Vogel concedes that close cousins to Dretske's example (such as the Car Theft Case) may undermine closure. The problem is that such examples appear to have features that cannot be exploited by arguments for (global) skepticism. Another supposed counterexample to closure has been presented by Robert Audi (1988, p. 77). For a critical discussion of Audi's argument see Feldman (1995). Mark Heller (1999) proposes an "Expanded Relevant Alternatives Principle" that, as he argues, avoids the problems of Dretske's original account. Peter Klein (1995) thinks that Dretske's argument against closure is correct regarding "externally situated evidence," but that it fails to realize that closure can be defended on internalist grounds. The reason, he says, is that the principle "does not require that the source of justification for the entailed proposition is anything other than the entailing proposition" (p. 221). My argument in this paper shows that if the notion of "externally situated evidence" is spelled out in terms of a Dretske-style theory of information, his argument does not even work for external evidence.

- ⁵ For groundbreaking work in the area see for example Cohen (1988), DeRose (1995), and Lewis (1996).
- ⁶ For a clear and representative statement of this approach see Cohen (2000, p. 103).
- ⁷ Sometimes Dretske is most explicit about this: “The *only* way to preserve knowledge of homely truths, the truths everyone takes themselves to know, is [...] to abandon closure” (Dretske forthcoming, my emphasis).
- ⁸ For a helpful discussion of these indirect arguments see for instance Williams (1996, pp. 330–336). Williams argues in chapter 8 of *Unnatural Doubts* that virtually all externalist attempts to show that closure fails are unsuccessful.
- ⁹ Put in the terminology Dretske introduces in *Knowledge and the Flow of Information*, we may say that, if *K* receives the information that *e*, and knows that *e* entails not-*h*, then the information that not-*h* is *analytically nested* in the information that *e*. Cf. Dretske (1981a, p. 71).
- ¹⁰ Cf. Dretske (1981a, p. 32): “Causality is a manifestation of a regular, lawlike, succession between events of type *C* and events of type *E* under relevantly similar conditions.”
- ¹¹ That Dretske’s theory falls victim to such counterexamples has been argued by Alvin Plantinga. See Plantinga (1993, pp. 195–197).
- ¹² For more detailed examinations of various epistemic closure principles see for example Brueckner (1985), Hales (1995), Luper (2001), or Barke (2002, pp. 26–43).
- ¹³ This is an embellishment considered for example by Brueckner (1985, p. 91). He argues that only a closure principle that has been altered along such lines will do justice to the skeptic, for “his target are not the careless epistemics” who fail to believe what they know follows from what they know. Instead, “the target knowers are ones who know that certain propositions are skeptical counterpossibilities to what they claim to know and who believe that these possibilities do not obtain” (1985, pp. 91f.).

REFERENCES

- Audi, R.: 1988, *Belief, Justification, and Knowledge*, Wadsworth, Belmont, CA.
- Barke, A. C. J.: 2002, *The Closure of Knowledge in Context*, Mentis, Paderborn.
- Brueckner, A.: 1985, ‘Skepticism and Epistemic Closure’, *Philosophical Topics* **13**, 89–117.
- Cohen, S.: 1988, ‘How to be a Fallibilist’, *Philosophical Perspectives* **2**, 91–123.
- Cohen, S.: 2000, ‘Contextualism and Skepticism’, *Philosophical Issues* **10**: *Skepticism*, edited by E. Sosa and E. Villanueva, 94–107.
- DeRose, K.: 1995, ‘Solving the Skeptical Problem’, *The Philosophical Review* **104**, 1–52. Reprinted in DeRose and Warfield 1999.
- DeRose, K. and T. Warfield (eds.): 1999, *Skepticism – A Contemporary Reader*, Oxford University Press, New York, Oxford.
- Dretske, F. I.: 1970, ‘Epistemic Operators’, *The Journal of Philosophy* **67**, 1002–1023. Reprinted in Dretske 2000, 30–47.
- Dretske, F. I.: 1971, ‘Conclusive Reasons’, *Australasian Journal of Philosophy* **49**, 1–22. Reprinted in Dretske 2000, 3–29.
- Dretske, F. I.: 1972, ‘Contrastive Statements’, *The Philosophical Review* **81**, 411–430.

- Dretske, F. I.: 1981a, *Knowledge and the Flow of Information*, Basil Blackwell, Oxford.
- Dretske, F. I.: 1981b, 'The Pragmatic Dimension of Knowledge', *Philosophical Issues* **40**, 363–378. Reprinted in Dretske 2000, 48–63.
- Dretske, F. I.: 1983, 'Précis of Knowledge and the Flow of Information', 'Open Peer Commentary', and 'Author's Response', *The Behavioral and Brain Sciences* **6**, 55–90.
- Dretske, F. I.: 1991, 'Dretske's Replies', in McLaughlin 1991, 180–221.
- Dretske, F. I.: 2000, *Perception, Knowledge, and Belief – Selected Essays*, Cambridge University Press, Cambridge.
- Dretske, F. I.: 2004, 'Externalism and Modest Contextualism', *Erkenntnis* **61**, 173–186.
- Dretske, F. I.: Forthcoming, 'The Case Against Closure', in E. Sosa and M. Steup (eds.), *Contemporary Debates in Epistemology*, Basil Blackwell, Oxford.
- Feldman, R.: 1995, 'In Defence of Closure', *Philosophical Quarterly* **45**, 487–494.
- Foley, R.: 1987, 'Dretske's 'Information Theoretic' Account of Knowledge', *Synthese* **70**, 159–184.
- Greco, J.: 2000, *Putting Skeptics in Their Place*, Cambridge University Press, Cambridge.
- Hales, St. D.: 1995, 'Epistemic Closure Principles', *The Southern Journal of Philosophy* **XXXIII**, 185–201.
- Heller, M.: 1999, 'Relevant Alternatives and Closure', *Australasian Journal of Philosophy* **77**, 196–208.
- Klein, P.: 1995, 'Skepticism and Closure: Why the Evil Genius Argument Fails', *Philosophical Topics* **23**, 213–236.
- Lewis, D.: 1996, 'Elusive Knowledge', *Australasian Journal of Philosophy* **74**, 549–567. Reprinted in DeRose and Warfield 1999.
- Luper, St.: 2001, 'Epistemic Closure Principles', *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/closure-epistemic/>.
- McLaughlin, B. P. (ed.): 1991, *Dretske and His Critics*, Basil Blackwell, Oxford.
- Plantinga, A.: 1993, *Warrant: The Current Debate*, Oxford University Press, Oxford.
- Vogel, J.: 1990, 'Are There Counterexamples to the Closure Principle?', in M. Roth and G. Ross (eds.), *Doubting: Contemporary Perspectives on Skepticism*, Kluwer Academic Publishers, Dordrecht.
- Williams, M.: 1996, *Unnatural Doubts – Epistemological Realism and the Basis of Skepticism*, Princeton University Press, Princeton.

Department of Philosophy
 University of Aberdeen
 Old Brewery, High Street
 Aberdeen, A24 3UB, UK
 E-mail: christoph.jaeger@abdn.ac.uk

MYLAN ENGEL, Jr.

WHAT'S WRONG WITH CONTEXTUALISM, AND A NONCONTEXTUALIST RESOLUTION OF THE SKEPTICAL PARADOX

ABSTRACT. Skeptics try to persuade us of our ignorance with arguments like the following: 1. I don't know that I am not a handless brain-in-a-vat [BIV]. 2. If I don't know that I am not a handless BIV, then I don't know that I have hands. Therefore, 3. I don't know that I have hands. The BIV argument is valid, its premises are intuitively compelling, and yet, its conclusion strikes us as absurd. Something has to go, but what? Contextualists contend that an adequate solution to the skeptical problem must: (i) retain epistemic closure, (ii) explain the intuitive force of skeptical arguments by explaining why their premises initially seem so compelling, and (iii) account for the truth of our commonsense judgment that we do possess lots of ordinary knowledge. Contextualists maintain that the key to such a solution is recognizing that the semantic standards for 'knows' vary from context to context such that in skeptical contexts the skeptic's premises are true and so is her conclusion; but in ordinary contexts, her conclusion is false and so is her first premise. Despite its initial attractiveness, the contextualist solution comes at a significant cost, for contextualism has many counterintuitive results. After presenting the contextualist solution, I identify a number of these costs. I then offer a noncontextualist solution that meets the adequacy constraint identified above, while avoiding the costs associated with contextualism. Hence, one of the principal reasons offered for adopting a contextualist theory of knowledge – its supposedly unique ability to adequately resolve the skeptical problem – is undermined.

The *intense* view of these manifold contradictions and imperfections in human reason has so wrought upon me, and heated my brain, that I am ready to reject all belief and reasoning, and can look upon no opinion even as more probable or likely than another. Where am I, or what? ...

Most fortunately it happens, that since reason is incapable of dispelling these clouds, nature herself suffices to that purpose, and cures me of this philosophical melancholy and delirium, either by relaxing this bent of mind, or by some avocation, and lively impression of my senses, which obliterate all these chimeras. I dine, I play a game of back-gammon. I converse, and am merry with my friends; and when after three or four hour's amusement, I would return to these speculations, they appear so cold, and strain'd, and ridiculous, that I cannot find in my heart to enter into them any farther (Hume, 1739, p. 268f).

There is an inconvenience which attends all abstruse reasoning, that it may silence, without convincing an antagonist, ... When we leave our closet, and engage in the common affairs of life, its conclusions seem to vanish, like the phantoms of the night on the appearance of the morning (Hume, 1739, p. 455).

David Hume



Erkenntnis 61: 203–231, 2004.

© 2004 Kluwer Academic Publishers. Printed in the Netherlands.

1. INTRODUCTION

Like the phantom-free Hume who has left his closet and is basking in the light of commonsense, I know a lot. I know that I have hands. I know that there is a laptop computer in front of me. I know that I am typing away on that computer with those very hands, *my* hands. I know that I am in funky coffeehouse. I know that the music of Nickel Creek is playing in the background. And, sadly, I know that my coffee cup is almost empty. These are just some of my many epistemic accomplishments. You know a lot, too. You know that you have eyes. You know that you are seeing a journal article (or a computer screen image of a journal article) with those eyes. You know that the article you are reading is written in English. Together, we know a lot. At least, we think we do, until we encounter the skeptic.

The skeptic contends that all of the above knowledge ascriptions are false. To convince us of our ignorance, the skeptic puts forth some skeptical hypothesis *H* such that if *H* were true, then (i) our evidential situation would be phenomenologically indistinguishable from our current evidential situation, and yet, (ii) all our commonsense perceptual beliefs would be false. The skeptic then argues that since we don't know that *H* is false, we don't know that the beliefs of commonsense are true. We might be being systematically deceived by the malevolent machinations of Descartes's demon or by the brain-distorting effects of Lehrer's Googols.¹ Or, consider the ever popular brain-in-a-vat [BIV] hypothesis, according to which I am a disembodied brain floating in a vat of nutrient hooked up to a sophisticated computer that is producing coherent experiences in me by stimulating my sensory cortex in exactly the same way it would have been stimulated were I perceiving normally. Since my being a *handless* BIV entails that I don't have hands, the skeptic argues as follows:

BIV arguments**BIV 1**

1. I don't know that I'm not a handless BIV.
2. If I don't know that I'm not a handless BIV, then I don't know that I have hands.
Therefore,
3. I don't know that I have hands.

Not wanting you to miss out on all the fun, the skeptic also argues as follows:

BIV 2

1. You don't know that you're not an eyeless BIV.
2. If you don't know that you're not an eyeless BIV, then you don't know that you have eyes.
Therefore,
3. You don't know that you have eyes.

These arguments are valid and their premises seem unassailable, and yet their conclusions strike us as absurd. Of course, accepting their premises while rejecting their conclusions commits us to what Stewart Cohen calls the inconsistent triad: 1, 2, and ~ 3 . Our options aren't appealing. Rejecting these arguments' second premises commits us to two rather unsavory abominable conjunctions: "I know that I have hands, but I don't know that I'm not a *handless* BIV" and "You know that you have eyes, but you don't know that you're not an *eyeless* BIV"; and rejecting these arguments' first premises requires our knowing that we are not BIVs, but how could we know that? And yet, surely, I know that I have hands (with which I am currently typing), and you know that you have eyes (with which you are currently reading!). We're in the skeptical fly bottle, alright. Something has to go, but what?

That's where contextualism comes in. In true Wittgensteinian fashion, contextualists promise to show us the way out of the fly bottle. But, they stress, not just any way out of the bottle will do: One can't just reject a premise and release the flies; a *satisfactory* solution to the skeptical problem must explain why we feel ensnared in the first place. As Keith DeRose puts it:

In seeking a solution to this puzzle, we should seek an explanation of how we fell into this trap in the first place, and not settle for making a simple choice among three distasteful ways out of the trap. We must explain how two premises that together yield a conclusion we find so incredible can themselves seem so plausible to us (DeRose, 1995, p. 3).²

In particular, contextualists insist that an adequate solution to the skeptical problem must meet the following three *desiderata*: (i) It must retain epistemic closure,³ (ii) it must explain the intuitive force of skeptical arguments by explaining why their premises initially seem so compelling, and (iii) it must account for the truth of our commonsense judgment that we do in fact possess lots of ordinary knowledge. The key to such a solution, according to contextualism, is recognizing that the semantic standards for 'knows' vary from context to context such that in skeptical contexts the skeptic's

premises are true and so is her conclusion; but in ordinary contexts, her conclusion is false and so is her first premise. Of course, there are embellishments. In what follows, I will provide those embellishments, for the devil is in the details. I will first explain the main tenets – the defining core – of the position that has come to be known as *epistemic contextualism*. Epistemic contextualism is viewed by some as a kind of panacea for nearly all of our epistemic ills. At the very least, it has been invoked to cure several persistent epistemological hangovers, including the lottery paradox and the Gettier problem, but its single biggest selling point by far is the apparent ease with which it can resolve the skeptical paradox. If it fails to deliver in this latter regard or if there is a preferable noncontextualist solution, then one of the primary motivations for contextualism collapses. In what follows, I present the contextualist “solution” to the skeptical problem. I then point out a number of counterintuitive costs associated with contextualism. If we can provide an alternative noncontextualist solution to the skeptical problem that meets the contextualist’s adequacy constraint on such solutions, while avoiding contextualism’s counterintuitive results, then the primary motivation for contextualism will have been undermined. In the paper’s final sections, I propose such a noncontextualist solution to the skeptical problem. Since epistemic contextualism is rooted in semantic contextualism, let’s begin our discussion there.

2. SEMANTIC CONTEXTUALISM

In “Scorekeeping in a Language Game”, David Lewis rightly observes that when interpreting people’s utterances, we employ various context-sensitive *rules of accommodation* that facilitate effective communication (Lewis, 1979). Lewis also stresses that contexts can shift very quickly, and we routinely succeed in accommodating these shifts. Here’s an embellished version of his famous cat case. We are sitting in a garden in Mainz, Germany drinking a beer, and you’ve just set your beer on the park bench next to you. I’m telling you about one of my cats back in Illinois, and I say: “The cat loves to chase string. When she judges that it’s been too long since we’ve played string, the cat picks up the string in her mouth and brings it over to me with a ‘Play-with-me-dammit!’ look in her eye. Look out, the cat is about to spill your beer!” You immediately interpret my latter use of ‘the cat’ to refer to the stray cat that has just jumped up on the bench next to you. Your beer safely in hand, I continue: “I’m sure when I return home,

the cat will greet me at the door, string in mouth.” You accommodate each of my uses of ‘the cat’. You realize that my first two uses refer to my cat in Illinois, my third use refers to the stray cat on the bench, and my fourth use once again refers to my cat back home. So far so good. Semantic contextualism is a truism about language.

In the case just described, speaker intentions [my intentions] determined which cat I was referring to when I used the expression ‘the cat’. Context did not determine which cat I was referring to, *I* did. To be sure, context played a role in allowing you to ascertain which cat I was referring to. Since a cat in Illinois is unlikely to spill a beer in Mainz, you rightly interpreted me to be talking about the stray cat in Mainz, and since a cat in Mainz won’t likely greet me when I return to Illinois, you rightly interpreted me to be talking about my own cat again. Epistemic contextualism takes as its starting point the truism that language is highly context-sensitive in just the way that semantic contextualists maintain.

3. EPISTEMIC CONTEXTUALISM

While epistemic contextualists differ on the details of their theories, e.g. some work within the relevant alternatives framework (see Cohen 1988) and others embrace the subjunctive conditionals approach to knowledge (see DeRose, 1995), there are certain core defining features on which all epistemic contextualists agree.

3.1. *The Metalinguistic Turn*

First, contextualists maintain that there is no correct context-independent standard of knowledge. Consequently, there is no context-independent fact of the matter as to whether or not *S* knows that *p*. Since there is no fact of the matter as to whether or not *S* knows that *p* outside a context of ascription, they maintain that we epistemologists should drop all talk about whether or not *S* knows that *p*. Our focus, instead, should be on whether sentences of the form ‘*S* knows that *p*’ are true in some specified context of ascription.

3.2. *Ascriber-Sensitive Contextualism*

The second defining feature of epistemic contextualism is that it is the context of the *knowledge-ascriber*, not the context of the

knowledge-ascriber, that determines which standards of knowledge are operant in her assertion. The truth conditions for ascriber *A*'s assertion '*S* knows that *p*' are determined by *A*'s context, not by *S*'s context. Consequently, for a given cognitive subject *S*, proposition *p*, and time *t*, it is possible for an ascriber *A1* to truthfully assert '*S* knows that *p*' at *t*, while another ascriber *A2* truthfully asserts '*S* does not know that *p*' at *t*, provided, e.g., that *A1* is in a low standards context and *A2* is in a high standards context. Given the different contexts of utterance, the proposition expressed by *A2*'s utterance is *not* the negation of the proposition expressed by *A1*'s utterance. So, *A1* and *A2* need not be disagreeing.

While the standard form of epistemic contextualism is *ascriber-sensitive contextualism*,⁴ as outlined above, it must be stressed that – unlike Lewis's cat example above – it's not up to the speaker/ascriber to decide what sense of 'knows' she is using. If the ascriber finds herself in a skeptical context – perhaps she's just been discussing deceptive demons in an epistemology class and is still worrying about them – the word 'knows' *in her mouth* will refer to high standards knowledge whether she intends it to or not. Even if she wants to use 'knows' in its low standards sense, she won't be able to as long as she is contemplating demons. It is the *context* of the ascriber, not the ascriber herself, that determines what standards of knowledge are operant in her assertion. As we shall subsequently see, the context of an ascriber *A* can be impacted and altered by the content of proposition *p* in *A*'s assertion '*S* knows that *p*'.

3.3. *Asymmetrical Standards Adjustments*

Epistemic contextualism differs from semantic contextualism in another important respect, as well. With Lewis's cat example, it was relatively easy to shift from Illinois cat to Mainz cat, and it was equally easy to shift from Mainz cat back to Illinois cat. But in the case of epistemic contextualism, there is an asymmetry in the amount of ease with which one can move from one use of 'knows' to another. Epistemic contextualists maintain that it is easy to move from a low standards context of 'knows' to a high standards context of 'knows', but they insist that it is very difficult to move from a high standards context back down to a low standards context. Once the epistemic standards have been raised, they tend to stay raised.⁵

4. CONTEXTUALISM TO THE RESCUE: SOLVING THE SKEPTICAL PROBLEM

We are now in a position to see how epistemic contextualism purports to solve the skeptical problem. As we go about our day moving from one mundane task to the next, I know that I have hands and you know that you have eyes. Better to state it metalinguistically: In an ordinary context, my knowledge self-ascription 'I know that I have hands' is true, and so is my knowledge attribution 'You know that you have eyes', because in an ordinary context, the operant standards of knowledge tend to be low. [In such an ordinary context, your knowledge self-ascription asserting that you know you have eyes (i.e. 'I know that I have eyes' *in your mouth*) would be true for similar reasons.]

Then, we meet a skeptic who presents us with the aforementioned BIV arguments. Prior to our encounter with the skeptic, the BIV hypothesis was not a relevant alternative, and hence, it was perfectly proper of us to ignore it. However, by the very act of mentioning the BIV hypothesis, the skeptic makes the BIV alternative relevant and improperly ignored, thereby in effect raising the standards of knowledge on us. Because I can't rule out the BIV alternative, the statement 'I don't know I'm not a BIV' is true *in my mouth* in that context, and as long as the BIV hypothesis remains relevant for me, 'I don't know that I have hands' is also true for me in that context. Likewise for you. Because you can't rule out the BIV hypothesis either, the sentence 'I don't know that I'm a BIV' is true *in your mouth* as well, and as long as it remains true for you, the sentence 'I don't know I have eyes' is also true *in your mouth*. That, according to contextualists, is why we feel so threatened by skeptical arguments. By presenting us with such arguments, the skeptic manipulates the semantic standards of 'knows' and moves us from an ordinary context in which low standards of knowledge are in effect to a skeptical context in which high standards are in effect; and we realize that, given the newly operant high standards of knowledge, we don't know the things we ordinarily take ourselves to know.

But, contextualists insist, the skeptic doesn't win the day, because like Hume once we leave our skeptical closet, return to our day-to-day affairs, and forget about the skeptic's challenge, our ordinary low standards become operant again, and our knowledge claims in that context are true. In fact, the skeptic was never able to show that we lacked low standards ordinary knowledge, because by presenting the skeptical argument the skeptic changed the topic from ordinary knowledge, where demons and BIV-scenarios are properly ignored, to

high standards knowledge, where these alternatives can't be ignored. And so, contextualists maintain, we have a solution to the skeptical puzzle that meets the contextualist's three stated *desiderata*. By the very act of asserting the BIV argument, the skeptic makes salient a skeptical alternative that we cannot rule out. In that context, the skeptic's premises are true, and so is her conclusion. That explains why we find the skeptic's argument so troubling. But when we leave our skeptical closets and forget about the skeptic's alternatives, these alternatives are no longer salient, and our ordinary low standards once again become operant, thus explaining why we remain convinced that we know that we have hands and eyes. Closure holds in any given context.

5. THE PRICE OF SUCCESS/THE COST OF CONTEXTUALISM

Despite the apparent ease with which it handles the skeptical problem, epistemic contextualism has numerous counterintuitive consequences. Here are a few of its unsavory results:

5.1. *Hume*

Contextualism entails that *from Hume's own perspective*, relative to the standards operant for knowledge self-ascriptions, Hume "knew" more in the bar when he was three sheets to the wind, than he did when he was carefully reflecting in his closet. Lewis makes the point as follows: "... when we do epistemology, and we attend to the proper ignoring of possibilities, we make knowledge vanish. First we do know, then we do not" (Lewis, 1996, p. 566). But that's not quite right. Strictly speaking, contextualism does not entail that any low standard knowledge is lost. Hume does not lose his low standard knowledge when he is entertaining skeptical reflections in his study, because *some other person* in a low standard context could still truthfully ascribe low standards knowledge to Hume, when Hume is in his closet.⁶ But I am not denying that there is *some perspective* from which one can truthfully attribute knowledge to Hume. The present worry is that *Hume* couldn't ascribe such knowledge to himself when he is in his skeptical closet. Contextualism implies that, *from Hume's own perspective*, he can't truthfully ascribe knowledge to himself when he is entertaining skeptical thoughts in his closet, but he can truthfully ascribe knowledge to himself when he is cavorting in the bar. Why? Because when Hume was contemplating various skeptical

hypotheses in his study, these alternatives were salient for him and could not properly be ignored. Of course, when he was making merry in the bar and ladies, not demons, were on his mind, evil demon scenarios were no longer relevant, and so, his slurred knowledge self-ascriptions were true. This is the sense in which contextualism entails that, *from Hume's own perspective*, he knew more in the bar when he was semi-inebriated than he did when he wrote *A Treatise of Human Nature* while carefully reflecting in his study. Contextualism also entails that when Hume was in the high standards context of his skeptical closet, he could not know that he still had low standards knowledge.⁷ One wonders how a person could write such an important and influential philosophical treatise, while, *from his own perspective*, knowing absolutely nothing (i.e. while, *from his own perspective*, all of his knowledge self-ascriptions are false). It strikes me as implausible in the extreme to think that more of Hume's knowledge self-ascriptions were true when he was half drunk in the bar than were true when he was soberly reflecting as carefully as possible on his epistemic situation. It also strikes me as implausible that Hume could have written a work of such lasting philosophical importance, while *from his own perspective* knowing absolutely nothing at all. How did he even manage to find a quill pen in his closet of absolute ignorance?

5.2. *Therapy 1*

I often find that when I go to philosophy conferences, esp. epistemology conferences, most of the participants know a lot more than I do. Not wanting to feel inferior, whenever I encounter these daunting intellectuals, I just contemplate Descartes's evil demon, and just like that, I can *truthfully* assert: 'These people don't know more than I. In fact, they don't know anything at all'. I feel better. Contextualist therapy at work.

5.3. *Therapy 2*

Sometimes I meet a particularly smug philosopher in need of therapy himself. He has not so subtly made it clear that he takes himself to know far more about the subject at hand than I do. In light of my recent therapy session above, I can truly assert 'He doesn't know anything at all', and I want to enlighten him to this truth. So, I walk up to him and say: 'You might be a BIV or a victim of wholesale

demonic deception. Deal with your total ignorance'. Just like that, I render all of his knowledge self-ascriptions false. If he claims to know any proposition in his newly acquired demon-context, his claim will be false. Now we're epistemic equals, and once again, I feel good.

But surely something is amiss. One can't make the sentence 'No one knows anything' true as easily as Therapy 1 suggests. Nor can one convince people of the truth of sentences ascribing a total state of ignorance to them, as simply Therapy 2 suggests.

5.4. *Unspeakable knowledge*

The contextualist is also committed to the view that there are many true propositions that we *know* are true, but that are such that it is impossible for us to truthfully assert that we know them. Consider the true proposition *that I am not a BIV*. According to our contextualist-sanctioned ordinary low standard knowledge, I know that this proposition is true, but I cannot truthfully assert 'I know that I am not a BIV', because in the very act of making such an assertion, I inevitably raise the standards for knowledge to a level so high that I no longer satisfy those standards with respect to the proposition *that I am not a BIV*. Such knowledge, which we do in fact possess, is not only unspeakable, it is unthinkable! As soon as we think about our not being BIVs, we cease to know we're not BIVs.⁸

5.5. *It's Hard to be a Consistent Contextualist*

If there ever were a skeptical context, this is it. I've been discussing deceptive demons, brain-distorting Googols, and BIV-scenarios, and my doing so has forced you to contemplate these skeptical alternatives. Before I mentioned these alternatives, they were irrelevant and properly ignored. But now that we're thinking about them, they are relevant and improperly ignored. So, if contextualism is correct, right now *in the present context*, my ascription 'You don't know that you have eyes' is true, and if you protest 'But I do know that I have eyes', your assertion is false. Since we are in a skeptical context, *in your mouth*, the utterance 'I know that I have eyes' is false. Moreover, not only is the utterance false, so is the thought. If right now you are thinking, 'How silly! Of course, I know I have eyes!', your belief is false, according to contextualism. By mentioning demons, Googols, and hallucination-producing vats, I've raised the standards for knowing, and you can't just lower them again at will. It takes awhile to forget what has been said, and until you do, *from your own perspective*,

you won't know that you have eyes [i.e. the knowledge ascription 'I know that I have eyes' *in your mouth* will be false]. But, I daresay, there is not a single person *reading* this article who doubts that she knows she has eyes. Not even the contextualists among us would think that the sentence 'I know I have eyes' is false *in their mouths/minds* while they are reading this article, despite my having just mentioned the BIV hypothesis. Despite the fact that contextualism entails that, in the present demon-infested context, no one reading this article knows that s/he has eyes, everyone reading it, including contextualists, continues to believe that s/he knows that s/he has eyes.

5.6. *Vanishing Force*

There is one phenomenon in particular that contextualism *in principle* cannot explain, namely, the widely recognized phenomenon that skeptical arguments tend to lose their force as we contemplate them time and again. Contextualism predicts just the opposite. It predicts that every time we contemplate a BIV argument, a demon argument, or any other skeptical argument, we will once again find ourselves in the throws of skepticism. It predicts that whenever we are presented with such an argument, we will conclude 'I really don't know what I thought I knew', but we conclude no such thing. The contextualist can't explain why the first time people encounter an argument like the BIV argument, they are often filled with epistemic angst, but as they study more epistemology, they cease to be moved by such arguments. To its credit, my noncontextualist solution does account for the vanishing force of skeptical arguments.

6. TOWARDS A NONCONTEXTUALIST RESOLUTION OF THE SKEPTICAL PROBLEM

I accept the contextualist's challenge that an adequate solution to the skeptical problem must do three things: (i) It must retain epistemic closure, (ii) it must explain the intuitive force of skeptical arguments by explaining why their premises initially seems so compelling, and (iii) it must account for the truth of our commonsense judgment that we do in fact possess lots of ordinary knowledge. Since I'm giving a noncontextualist solution to the skeptical problem that retains epistemic closure while maintaining that we do possess a great deal of knowledge, I must attack premise 1, but in the course of doing so, I must also explain why premise 1 *initially* strikes us as being so

plausible. Often when presented with a paradoxical argument like the BIV argument, the action is really taking place at the level of unstated assumptions. Why is it that we are initially so convinced by the skeptic's first premise? How does the skeptic persuade us to accept that premise? As already noted, she does so by presenting us with a detailed skeptical hypothesis *H* such that if *H* were true, then (i) our evidential situation would be phenomenologically indistinguishable from our current evidential situation, and yet, (ii) all our common-sense perceptual beliefs would be false. She then argues that the mere *possibility* of *H* is sufficient to prevent us from having any perceptual knowledge at all. It is important that *H* be possible. No one was ever moved to the brink of skepticism by the skeptical hypothesis: "You might be being deceived by a malevolent roundsquare". In presenting her skeptical hypothesis, the skeptic asserts that it is *possible* that I am a BIV. Starting with this premise, she then argues in one of two ways. Either she uses the *Argument from Possibility 1* to defend the first premise of BIV1 as follows:

Argument from Possibility 1 (AP1)

1. It is possible that I am a handless BIV.
2. If it's possible that I am a handless BIV, then I don't know that I'm not a handless BIV.
Therefore,
3. I don't know that I'm not a handless BIV.

Or, she uses the *Argument from Possibility 2* to argue straight away for the conclusion that I don't know that I have hands:

Argument from Possibility 2 (AP2)

1. It is possible that I am a handless BIV.
2. If it's possible that I am a handless BIV, then it's possible that I don't have hands.
3. If it's possible that I don't have hands, then I don't know that I do have hands.
Therefore,
4. I don't know that I have hands.⁹

Like its BIV1 counterpart, AP1 looks valid and its premises certainly seem to be true. So, we seem to have a good argument for AP1's conclusion, an argument which makes premise 1 of BIV1 *prima facie* plausible. AP2 also appears valid by repeated instances of *modus ponens*, and its premises likewise seem true. So, AP2 seems

to provide compelling reason to think that I don't know that I have hands. In what follows, I will argue that AP2 either begs the question or equivocates, and either way, fails to give us a good reason to accept AP2's skeptical conclusion. Since similar arguments *mutatis mutandis* will show that AP1 also either begs the question or equivocates, and so, fails to provide a good reason for accepting premise 1 of BIV1, I won't rehearse those arguments. To set the stage for my argument, I must briefly discuss epistemic possibility.

6.1. *Epistemic Possibility: Two Cases*

Consider the following case:

Philosopher Bob

Ordinary Joe and Philosopher Bob are sitting along the lakeshore in Chicago discussing Jim Java's whereabouts.

Joe: I haven't seen Jim Java in a few days. Perhaps, he is in New Orleans. He told me he was going to take a trip there soon, just so he could drink chicory coffee in the French Quarter.

Bob: It's not possible that Jim Java is in New Orleans. I just saw him at the Bump and Grind Coffeehouse twenty minutes ago.

Bob intends to assert something true. But on a metaphysical reading, his modal claim is false,¹⁰ because there is a μ -possible [metaphysically possible] world where Jim was at the Bump and Grind twenty minutes ago, but also where intercity transportation is remarkably efficient, so efficient in fact that one can get from Chicago to New Orleans in ten minutes. Of course, Bob knows that that world is not actual. In the actual world, it takes twenty minutes just to get a cab in Chicago, another thirty minutes to get to O'Hare Airport, and an hour to get through security. Given what Bob knows, Jim simply couldn't be in New Orleans. He couldn't have even made it to O'Hare in so little time. Bob's background knowledge entails that Jim is not in New Orleans, and Bob recognizes that entailment. Given Bob's knowledge, it is not e-possible [epistemically possible] for Jim to be in New Orleans. That is why Bob asserts what he does. Bob is making a true epistemic modal claim, not a false metaphysical one.

Consider a second case due to Kripke:

Goldbach's Conjecture

Goldbach's Conjecture (GC) is the mathematical conjecture that every even number greater than 2 is the sum of two primes. While no counter-instance to GC has ever been found, no one has ever demonstrated that GC is true. Since GC is a mathematical proposition, it has its truth-value necessarily. So, *if* GC is true (as is widely believed), then it is *necessarily* true; and yet, since it has never been proven, it seems true to say that it *might* be false. Such is the nature of conjectures. They might be false. How *might* a necessarily true proposition be false? In what sense, is it *possible* for a necessarily true proposition to be false? Kripke's answer is, "in the epistemic sense." He rightly observes that the 'might' and the 'possible' are being used in an epistemic sense merely to express our present ignorance of the truth value of GC (Kripke 1980, pp. 36–38).

6.2. A Stipulative Account of E-Possibility

As Ian Hacking has observed, certain "occurrences of *possible* can be modified by many adverbs of the form ϕ -ly: technically, economically, theoretically, medically, metaphysically, humanly" (Hacking 1975, p. 325). We can understand ϕ -possibility ascriptions using the following ϕ -possibility schema:

$\Diamond_{\phi} p$ p is ϕ -ly possible for S iff nothing of a ϕ -al sort ϕ -ly **precludes** the truth of p .

According to this schema, p is logically possible for S iff nothing of a logical sort *logically precludes* the truth of p ; and p is physically possible iff nothing of a physical sort *physically precludes* that p . Similarly, p is epistemically possible for S iff nothing of an epistemic sort *epistemically precludes* the truth of p . Let us, therefore, turn our attention to the notion of "epistemic preclusion" since it holds the key to understanding epistemic possibility.

6.3. Epistemic Preclusion

There are two ways that something known by S can *epistemically preclude* the truth of p for S . The first and most obvious way for p to be epistemically precluded for S is for S to know that $\sim p$. Let's call this "direct e-preclusion" and define it as follows:

D₁ p is *directly e-precluded* for S at t iff S knows that $\sim p$ at t .

To see that there is another way for something S knows to epistemically preclude p for S consider the following scenario.

Simone

I am lying awake in bed. Simone, one of my cats, is curled up against my side purring loudly. I know that Simone is in the bedroom, but I do not know that she is not in the kitchen, because I haven't bothered to form the belief that she is not in the kitchen. Still, my knowledge that she is in the bedroom, together with my background knowledge, self-evidently entails that she is not in the kitchen, i.e. this entailment is one I would immediately recognize, were I to consider it. My knowledge that Simone is in the bedroom indirectly e-precludes her being in the kitchen. Call this *indirect e-preclusion*:

D₂ p is *indirectly e-precluded* for S at t iff (i) S does not know that $\sim p$ at t , but (ii) S could come to know that $\sim p$ at t , strictly on the basis of propositions S knows at t .¹¹

With these definitions in hand, we can define *e-possibility* as follows:

D₃ p is e-possible for S iff p is neither directly nor indirectly e-precluded for S .

Or equivalently:

D'₃ p is e-possible for S at t iff (i) S does not know that $\sim p$ at t , and (ii) S could not come to know that $\sim p$ at t , strictly on the basis of propositions S knows at t .^{12,13}

D'₃ yields the right results where Philosopher Bob is concerned, for it entails that it's *not* e-possible for Bob that *Joe is in New Orleans*, because Bob knows that (i) Joe was at the Bump and Grind Coffeehouse in Chicago twenty minutes earlier and he also knows that (ii) in the actual world it takes more than twenty minutes to get from Chicago to New Orleans; and Bob recognizes that (i) and (ii) entail that Joe is not in New Orleans.¹⁴ D'₃ also yields the right results where Goldbach's conjecture is concerned. I do not know that GC is true, and nothing I currently know self-evidently entails that GC is true. [Neither I nor any expert mathematician has been able to see how to derive GC.] So, D'₃ entails that $\sim$ GC is e-possible for me and *ipso facto* that it is e-possible for me that GC is false.

6.4. The “Possibility” of the BIV Hypothesis

One problem with the *Argument from Possibility* (AP) should be obvious – it does not specify the kind of possibility it is employing in its premises. To help us sort through the various readings of AP2, let’s symbolize it as follows:

- AP2 1. $\Diamond b$
 2. $\Diamond b \rightarrow \Diamond \sim h$
 3. $\Diamond \sim h \rightarrow \sim Kh$
 $\therefore$ 4. $\sim Kh^{15}$

Perhaps, the skeptic hopes to establish the skeptical conclusion by appealing to the mere metaphysical possibility of my being a BIV as follows:

- MAP2 1. $\Diamond_{\mu} b$
 2. $\Diamond_{\mu} b \rightarrow \Diamond_{\mu} \sim h$
 3. $\Diamond_{\mu} \sim h \rightarrow \sim Kh$
 $\therefore$ 4. $\sim Kh^{16}$

MAP2 is valid, but it is not sound. Premise 3 is false. We cannot, generally speaking, derive epistemic conclusions from purely metaphysical premises.¹⁷ And we certainly cannot derive *S*’s ignorance of *p* from the mere μ -possibility of $\sim p$, for, as we have already seen with Philosopher Bob, the μ -possibility of $\sim p$ does not imply that *S* does not know that *p*. The μ -possibility that Jim Java is in New Orleans does not imply that Philosopher Bob does not know that Jim is not in New Orleans. Bob does know that Jim is not in New Orleans, even though, being a philosopher, he also knows that it is μ -possible that Jim is there.¹⁸

If the skeptic wishes to derive the epistemic conclusion that I don’t know I’m not a BIV from AP2, then she must appeal to epistemic possibility throughout her argument:

- EAP2 1. $\Diamond_e b$
 2. $\Diamond_e b \rightarrow \Diamond_e \sim h$
 3. $\Diamond_e \sim h \rightarrow \sim Kh$
 $\therefore$ 4. $\sim Kh^{19}$

6.5. *Exploring the Epistemic Possibility of the BIV Hypothesis*

Unlike MAP2, EAP2 presents us with a genuine skeptical problem – a paradox of its own, for it looks valid and its premises seem to be true, and yet, its conclusion strikes most epistemologists as false.²⁰ That premise 3 is true can be demonstrated as follows: D'_3 entails the following necessary condition for e-possibility: p is e-possible for S *only if* S does not know that $\sim p$.²¹ Formally, the condition can be stated as follows: $(p)(\Diamond_e p \rightarrow \sim K \sim p)$.²² Since premise 3 is an instance of this general truth, premise 3 is clearly true. Premise 2 is also true. The e-possibility of my being a handless BIV does imply the e-possibility of my having no hands; for if nothing I know e-precludes my being a handless BIV, then nothing I know e-precludes my having no hands. Given EAP2's validity and the truth of its premises 2 and 3, it follows that if premise 1 of EAP2 is true, then EAP2's conclusion 4 must also be true. And premise 1 looks true – it certainly seems e-possible that I am a handless BIV. Of course, it also seems clear that I know that I have hands. After all, I have a reliably produced, perceptually justified true belief that I have hands, and there is no Gettier funny-business going on. The paradox generated by EAP2 then is this: The BIV hypothesis is e-possible, and yet we know things incompatible with its e-possibility. How can that be? In what follows, I will resolve the paradox in a way that allows us to retain our philosophical intuition that the BIV hypothesis is e-possible, while also allowing us to retain our commonsense intuition that we do have knowledge of the external world around us. I will argue that EAP2 fails to provide us with a reason to think that external world skepticism is true, because it either rests on an equivocation or it begs the question.

6.6. *Two Kinds of E-Possibility*

I submit that, even in EAP2, there lurks a hidden ambiguity, because there is more than one kind of e-possibility. Let me explain. One of the reasons it is difficult to get an intuitive handle on e-possibility is because our intuitive e-possibility assessments are split along the same infallibilistic/fallibilistic lines as our ordinary epistemic evaluations.²³ We typically relativize our e-possibility assignments to the propositions we *fallibly* know [know_f], but occasionally, we make our e-possibility assignments relative to the propositions we *infallibly* know [know_i].²⁴ To avoid conflating our fallibilistic e-possibility

assessments with our infallibilistic e-possibility assessments, D'_3 must be revised as follows:

- D'_{3f} p is e-possible_f for S at t iff (i) S does not know_f that $\sim p$ at t , and (ii) S could not come to know_f that $\sim p$ at t , strictly on the basis of propositions S knows_f at t .
- D'_{3i} p is e-possible_i for S at t iff (i) S does not know_i that $\sim p$ at t , and (ii) S could not come to know_i that $\sim p$ at t , strictly on the basis of propositions S knows_i at t .²⁵

Once we recognize the distinction between fallibilistic and infallibilistic e-possibility, EAP2 itself turns out to be multiply ambiguous between a purely infallibilistic reading [EAP2i], two mixed readings [EAP2_{x1} and EAP2_{x2}], and a purely fallibilistic reading [EAP2f]:

EAP2i	1. $\Diamond_{ei}b$	EAP2 _{x1}	1. $\Diamond_{ei}b$
	2. $\Diamond_{ei}b \rightarrow \Diamond_{ei}\sim h$		2. $\Diamond_{ef}b \rightarrow \Diamond_{ef}\sim h$
	3. $\Diamond_{ei}\sim h \rightarrow \sim K_ih$		3. $\Diamond_{ef}\sim h \rightarrow \sim K_fh$
$\therefore$	4. $\sim K_ih$ ²⁶	$\therefore$	4. $\sim K_fh$ ²⁷
EAP2 _{x2}	1. $\Diamond_{ei}b$	EAP2f	1. $\Diamond_{ef}b$
	2. $\Diamond_{ei}b \rightarrow \Diamond_{ef}\sim h$		2. $\Diamond_{ef}b \rightarrow \Diamond_{ef}\sim h$
	3. $\Diamond_{ef}\sim h \rightarrow \sim K_fh$		3. $\Diamond_{ef}\sim h \rightarrow \sim K_fh$
$\therefore$	4. $\sim K_fh$	$\therefore$	4. $\sim K_fh$

Neither mixed reading is plausible. EAP2_{x1} identifies an equivocal reading of EAP2 and is simply invalid. As for EAP2_{x2}, its second premise is false. The fact that none of the propositions I know_i e-precludes that I am a BIV does not entail that none of the propositions I know_f e-precludes my not having hands. Why? Because I might know_f that *I have hands* – and that knowledge_f would fallibilistically directly e-preclude that I don't have hands – without knowing_i that I am not a BIV and without knowing_i anything that entails or justifies_i me in believing that I'm not a BIV.

As for EAP2i, while it is clearly sound, it isn't of much philosophical interest. First to its soundness. EAP2i is valid, and its second and third premises are true for reasons analogous *mutatis mutandis* to those offered when discussing premises 2 and 3 of EAP2 above. That leaves premise 1 of EAP2i to consider. It is generally acknowledged that we have very little in the way of knowledge_i, for our evidence rarely entails that for which it is evidence. We may know_i a few *cogito*

propositions, but that's about it. Given the little, if anything, that we know_i, very few propositions, if any, are infallibilistically e-precluded for us. Suppose, for example, that I possess no knowledge_i whatsoever. Then, no propositions are infallibilistically e-precluded for me, and so every proposition is e-possible_i for me. If, on the other hand, I do possess *cogito* knowledge_i of my own existence, then that knowledge_i infallibilistically e-precludes my own nonexistence for me. But my *cogito* knowledge_i does not infallibilistically e-preclude my being a BIV, because I cannot justifiably_i infer my nonenvattedness from the few *cogito* propositions I know_i. Since my being a BIV is not infallibilistically e-precluded for me, my being a BIV is e-possible_i for me, just as premise 1 of EAP2i asserts.²⁸ Hence, EAP2i is sound. I don't know_i that I have hands.

The reason EAP2i is uninteresting is because I don't need to contemplate the e-possibility_i of far-fetched BIV hypotheses to realize that I lack knowledge_i that I have hands. Presumably, my current visual and tactile experiences are what justify me in believing that I have hands, and it is obvious that those experiences *do not entail* that I have hands, for I can have those same experiences as a result of dreams, mushroom-induced hallucinations, virtual reality machines, phantom limb experiences, etc. Therefore, I am not justified_i in believing that I have hands. Since I lack justification_i for believing that I have hands, and since justification_i is necessary for knowledge_i, it follows that I lack knowledge_i that I have hands. So, we do not need to appeal to EAP2i to establish such a conclusion.

The only interesting version of EAP2 is EAP2f, since it is the only version which threatens to undermine our ordinary fallibilistic knowledge of the objects around us. Like EAP2i, EAP2f is valid and its second and third premises are true, again for reasons analogous *mutatis mutandis* to those offered in support of EAP2's premises 2 and 3. The problem with EAP2f is that the skeptic is in no position to assert that its first premise is true, for *suppose* that I know_f that I have hands (*h*). Then I know_f a proposition – namely, *h* – that entails that I am not a handless BIV.²⁹ And since I have the cognitive power needed to grasp this entailment, my knowing_f that *h* would fallibilistically (indirectly) e-preclude for me my being a handless BIV and would thus render the BIV hypothesis e-impossible_f for me (regardless whether I have actually noticed the entailment or not). So, the skeptic can only rationally assert EAP2f's first premise – that it is e-possible_f for me that I am a handless BIV – if she *assumes* the truth of EAP2f's conclusion. Since the skeptic cannot rationally assert EAP2f's first premise without assuming that I lack knowledge_f

that I have hands, she cannot assert premise 1 without assuming the very thing in question. Granted, if the skeptic could give an independent reason for thinking EAP2f's first premise true – a reason that did not make reference to the truth of EAP2f's conclusion, then EAP2f would not beg the question; but she can't because e-possibility_f is analyzed in terms of knowledge_f. Consequently, EAP2f effectively begs the question, because to be rationally entitled to assert premise 1 of EAP2f, the skeptic must *first* be rationally entitled to assert that $\sim K_f h$. Perhaps the skeptic can provide some other argument for $\sim K_f h$, which she can then use to establish $\sim K_f h$ and *ipso facto* EAP2f's first premise. But then, it is this other argument – not EAP2f – that is doing all the skeptical work. Any argument A1 such that one must first establish *the conclusion* of A1 via some second argument A2 *before* one can rationally assert the premises of A1 is itself worthless in establishing the conclusion of A1. EAP2f is such an argument. In order for the skeptic to rationally assert premise 1 of EAP2f, she must first prove the truth of EAP2f's conclusion with a different argument, thereby rendering EAP2f superfluous.³⁰

6.7. *Undermining AP2 and Resolving the EAP2 Paradox*

An adequate solution to the skeptical problem must not only explain where the skeptic's argument goes wrong, it must also explain why the skeptic's argument initially has such strong intuitive appeal. My solution does both. The argument from possibility AP2 goes wrong, because it is either unsound (due to a false premise as in MAP2 and EAP2_{x2}), or uninteresting and irrelevant to *fallible* knowledge (as in EAP2i), or invalid (due to equivocation as in EAP2_{x1}), or question-begging (as in EAP2f). As a result, AP2 provides no good reason for thinking that I lack knowledge_f that I have hands. Why then are so many people caught in AP2's skeptical grip, when first presented with the argument? The answer is that either: (i) having initially been drawn in by the μ -possibility of the BIV hypothesis, they conflate μ -possibility with e-possibility, thereby, in effect, illegitimately drawing an epistemic conclusion from purely metaphysical premises, or (ii) they recognize that the argument must be couched in terms of e-possibility, but they fail to notice the subtle equivocation between fallibilist and infallibilist senses of epistemic possibility identified in EAP2_{x1}. Given the subtlety of each mistake, it is perfectly understandable that one find the unqualified AP2 initially threatening, indeed.

This way of undermining the skeptical problem posed by AP2 also allows us to resolve the paradox that EAP2 itself generates. The reason we are inclined to think that the BIV hypothesis *is* e-possible even though we *know* things incompatible with its e-possibility is because in making our BIV e-possibility assessment we are making an e-possibility_i assessment, whereas in claiming to know that I have hands, we're making a knowledge_f claim. Since we have very little, if any, *infallible* knowledge, nothing we *infallibly* know e-precludes the truth of the BIV hypothesis, and so, the BIV hypothesis is e-possible_i for us. That is why we are initially seduced into accepting premise 1 of the ambiguous EAP2. The epistemological mistake that has been repeated for centuries and that most people make when first confronted with the e-possibility_i of the BIV scenario is concluding, on that basis, that we have no knowledge_f, which is just to fall prey to the equivocation identified in EAP2_{X1}. The e-possibility_i of the BIV hypothesis does prevent us from having knowledge_i of the existence of the external world, as EAP2_i shows. That's as it should be. But EAP2 is impotent when it comes to knowledge_f, because there is no nonquestion-begging way to establish the e-possibility_f of the BIV hypothesis.

This solution has the added virtue that it can explain why skeptical arguments typically lose their force for epistemologists. Once we realize that the skeptic cannot assert her major premise – that it's e-possible_f that we are BIVs – without begging the question, we no longer find her arguments compelling. At first, we don't realize that she cannot rationally assert that premise without assuming the truth of her conclusion, but once we do, we find her arguments to be of absolutely no use in defending skepticism. This, I suspect, is why many epistemologists are no longer bothered by the very same skeptical arguments that at one time troubled them deeply.

7. EPILOGUE: BACK TO THE VAT

At this point some of you may feel cheated. Haven't I pulled an end run around the skeptic (rather than confronting her) by substituting the argument from possibility for the BIV argument and arguing that the former is question-begging? I can imagine a skeptic objecting as follows:

You may be right that the argument from possibility begs the question, but what about the BIV argument with which we began? It seems immune to your question-begging charge, since it's valid, its second premise BIV1.2 follows from the closure

principle, and there is an independent argument for its first premise BIV1.1, an argument that makes no reference to BIV1's conclusion [BIV1.3].

I submit that the standard argument offered in support of BIV1.1 does presuppose the truth of BIV1.3. To see why, consider the following dialogue with me playing the role of the nonskeptic:

Skeptic: Do you know_f that you have hands?

Me: Why yes, I think I do.

Skeptic: There's good reason to think you are mistaken.

Me: Really? I'd need a *very* good reason to give up such a commonsensical belief as the belief that I know_f that I have hands. So tell me, what's the reason?

Skeptic: It's a very good reason, alright. It's based on the extremely plausible principle of epistemic closure according to which: If *S* knows_f that *p*, and *S* knows_f that *p* entails *q*, and *S* considers *q* in light of her knowledge_f that *p* and that *p* entails *q*, and *S* bases her belief that *q* on this knowledge_f, then *S* knows_f that *q*. Surely, you accept the principle of epistemic closure, don't you?

Me: I certainly do.

Skeptic: Good, then I've got you. Here's the argument: (i) You don't know_f you're not a handless BIV. The second premise derives from the closure principle, which you've just embraced: (ii) If you don't know_f you're not a handless BIV, then you don't know_f that you have hands. Therefore, (iii) you don't know_f that you have hands.

Me: Hmm. Why should I accept (i)? It seems to me that you are presupposing that I don't know_f I have hands in your first premise.

Skeptic: No, I'm not. I wouldn't expect you to accept (i) without an independent argument supporting it, but I have such an argument. It's the famous argument from possibility and runs as follows: (iv) Surely, it's *possible* that you are a handless BIV. (v) If it's possible that you are a handless BIV, then you don't know_f that you're not a handless BIV. So, (vi) you don't know_f you're not a handless BIV – which is just my original premise (i).

Me: What kind of *possibility* are you appealing to in (iv)?

Skeptic: Epistemic possibility, of course.

Me: Fallibilistic or infallibilistic e-possibility?

Skeptic: Fallibilistic.

Me: Can we agree to the following account of fallibilistic e-possibility:

D'_{3f} p is e-possible_f for S at t iff (i) S does not know_f that $\sim p$ at t , and (ii) S could not come to know_f that $\sim p$ at t , strictly on the basis of propositions S knows_f at t ?

Skeptic: Yes, something like that seems correct.

Me: So, given D'_{3f} , your argument from possibility [our AP1] in support of (i) reduces to the following argument: (iv*) I don't know_f that I'm not a handless BIV, and I could not come to know_f that I'm not a handless BIV, strictly on the basis of the propositions I currently know_f. (v*) If I don't know_f that I'm not a handless BIV, and I could not come to know_f that I'm not a handless BIV, strictly on the basis of the propositions I currently know_f, then I don't know_f that I'm not a handless BIV. Therefore, (vi) I don't know_f that I'm not a handless BIV. Right?

Skeptic: Right.

Me: But then your argument from possibility fails to provide me with a nonquestion-begging reason to accept the first premise of your BIV argument. After all, the first conjunct of premise (iv*) of your argument from possibility is "I don't know_f that I'm not a handless BIV", which is precisely what your argument from possibility was supposed to establish.

Skeptic: Hmm.

Me: Worse still, you can only rationally assert the second conjunct of premise (iv*) – that I could not come to know_f that I am not a handless BIV, strictly on the basis of propositions I currently know_f – if you make certain presuppositions about what I currently know_f. In particular, you must assume that I don't currently know_f that I have hands; for given your own commitment to closure, *if* I do currently know_f that I have hands, then I could easily come to know_f that I am not a *handless* BIV on the basis of my knowledge_f that I have hands. Thus, in offering your argument from possibility in support of premise (i) of your BIV argument, you are presupposing the truth of the very conclusion your BIV argument was supposed to establish. You promised to provide a good argument for me to accept (iii) – the conclusion that I do not know_f that I have hands – and now it turns out that the only way for you to defend the first premise of your

original BIV argument is to appeal to another argument which presupposes the truth of (iii) in its first premise, i.e. premise (iv). Such a pair of arguments is of no use in proving the truth of conclusion (iii). If there is a good argument for the skeptical conclusion “I don’t know_f that I have hands”, you certainly have not produced it. Absent a good nonquestion-begging reason to believe that I don’t know_f that I have hands, it seems perfectly reasonable to retain the commonsense belief that I do know_f that I have hands and, given closure, it’s also reasonable to believe that I know_f that I’m not a handless BIV.

ACKNOWLEDGEMENTS

An earlier version of this paper was presented at the “Contextualist Approaches to Epistemology” Conference in Mainz, Germany. I wish to thank my commentator, Gilbert Scharifi, for his thoughtful commentary on my paper. I also wish to thank the conference participants for their helpful comments during the discussion period. Special thanks to Matthias Steup, Bruce Russell, John Carroll, Tomis Kapitan, James Dye, Elke Brendel, Christoph Jäger, and an anonymous referee for this journal for their extremely helpful, detailed comments on earlier drafts. The present version is much improved as a result.

NOTES

- ¹ Skeptical hypotheses need to be sufficiently detailed alternative explanations of our experiences to give us pause. Lehrer details the Googol hypothesis as follows: “There are a group of creatures in another galaxy, call them Googols, whose intellectual capacity is 10^{100} that of men, and who amuse themselves by sending out a peculiar kind of wave that affects our brain in such a way that our beliefs about the world are mostly incorrect. This form of error infects beliefs of every kind, but most of our beliefs, though erroneous, are nevertheless very nearly correct. This allows us to survive and manipulate our environments” (Lehrer, 1971, p. 356).
- ² Cohen endorses essentially the same adequacy constraint: “The burden of the fallibilist is to resolve these puzzles and paradoxes in a way that preserves the truth of our everyday knowledge attributions. But a *satisfying* resolution requires an explanation of why the paradox arises – an explanation of why we have the intuitions that saddle us with the paradox” (Cohen, 1988, p. 94).
- ³ Epistemic closure is the thesis that knowledge is closed under known logical implication. Of course, as is widely recognized, we must be careful in how we

formulate the closure principle. For example, overly simplistic formulations, like Nozick's $[K_s(p) \& K_s(p \rightarrow q)] \rightarrow K_s(q)$, are false, for S may simply fail to put two and two together and thus not come to believe that q , in which case S will fail to know that q . [After presenting the above formulation, Nozick himself adds: "this principle counts on the person to draw the inference to q " (Nozick, 1981, p. 205).] To avoid confusions that can be caused by such overly simplistic formulations, let us agree to understand the epistemic closure principle as follows: If S knows that p , and S knows that p entails q , and S considers q in light of her knowledge that p and that p entails q , and S bases her belief that q on this knowledge, then S knows that q .

⁴ DeRose refers to his version of contextualism as "attributor contextualism" (DeRose, 1999, p. 190).

⁵ Lewis makes the point as follows: "I take it that the rule of accommodation can go both ways. But for some reason raising of standards goes more smoothly than lowering. If the standards have been high, and something is said that is true enough only under lowered standards, and nobody objects, then indeed the standards are shifted down. But what is said, although true enough under the lowered standards, may still seem imperfectly acceptable. Raising of standards, on the other hand, manages to seem commendable even when we know that it interferes with our conversational purpose. Because of this asymmetry, a player of language games who is so inclined may get away with it if he tries to raise the standards of precision as high as possible – so high, perhaps, that no material object whatever is hexagonal" (Lewis, 1979, p. 352f). "We get the impression that the sceptic ... has the last word. Again this is because the rule of accommodation is not fully reversible. For some reason, I know not what, the boundary readily shifts outward if what is said requires it, but does not so readily shift inward if what is said requires that" (Lewis, 1979, p. 355).

⁶ See DeRose (2000) for a fuller discussion of "lost knowledge".

⁷ For a formal demonstration that, in high standards contexts, one cannot know that one has low standards knowledge, see Brendel (2003).

⁸ DeRose acknowledges this very point: "Thus, on our solution, we do know, for instance, that we're not BIVs, according to ordinary low standards for knowledge. But, though (1) [of our BIV argument] is false when evaluated according to those ordinary low standards, we're able to explain its plausibility, as we've seen, by means of the fact that the high standards at which (1) is true are precisely the standards that an assertion or denial of it put into play. Since attempts to assert (1) are bound to result in truth, and attempts to deny it are destined to produce falsehood, it's no surprise that we find it so plausible" (DeRose, 1995, p. 39f). DeRose is so eager to explain the plausibility of premise (1) of the BIV argument that he doesn't seem to notice how counterintuitive it is to maintain that people have all sorts of unspeakable and unthinkable knowledge.

⁹ Some philosophers regard *arguments from possibility* as the most fundamental skeptical arguments. See, e.g. Vogel (2005, 72), where he writes: "The argument which supports skepticism is one of the most famous in the history of philosophy. It turns on the possibility that we might be victims of some kind of massive sensory deception." Also see: (a) Stroud (1984, chapter 1), where he develops the Cartesian argument from possibility at length; and (b) Nozick's discussion of "Skeptical Possibilities" (Nozick, 1981, pp. 198–204).

- ¹⁰ For present purposes, let us stipulatively agree to the following: A proposition p is μ -possible *iff* there is a μ -possible world where p is true. A proposition p is μ -impossible *iff* there are no μ -possible worlds where p is true.
- ¹¹ Obviously, condition (ii) needs unpacking. Here is what is intended by condition (ii):
- D₄ S could come to know that $\sim p$ at t , strictly on the basis of the propositions S knows at t , *iff* either (1) one or more of the propositions S knows at t self-evidently entail that $\sim p$ for S (such that it is within S 's cognitive capacity at t to grasp that entailment immediately at t); or (2) $\sim p$ is true, one or more of the propositions S knows at t provide an adequate justificatory basis for believing that $\sim p$, and it is within S 's cognitive capacity at t to see that these propositions justify her in believing that $\sim p$.
- ¹² It's worth noting that D'₃ entails that no true proposition is ever e-impossible for S , because if p is true, then $\sim p$ is false, and since $\sim p$ is false, S couldn't know that $\sim p$. At first blush, it might seem surprising that no true proposition is ever e-impossible. Nevertheless, the result is correct. How, after all, could something S *knows* to be true e-preclude the truth of another true proposition?
- ¹³ D'₃ can be stated more explicitly as follows:
- D''₃ p is e-possible for S at t *iff* (i) S does *not* know that $\sim p$ at t ; (ii) if one or more of the propositions S knows at t entail that $\sim p$, then it is *not* within S 's cognitive capacity at t to grasp that entailment; and (iii) if $\sim p$ is true and if one or more of the propositions S knows at t provide an adequate justificatory basis for believing that $\sim p$, then it is *not* within S 's cognitive capacity at t to see that these propositions justify her in believing that $\sim p$.
- ¹⁴ D'₃ also entails the right result regarding **Simone**. I know that Simone is in the bedroom, because I reliably, justifiably, and truly believe that she is lying in bed next to me. I do not know that *Simone is not in the kitchen*, because I haven't considered that proposition. However, some of the propositions I know – (a) that *Simone is in the bedroom*, and (b) that *the bedroom and the kitchen are distinct rooms in my palatial estate* – self-evidently entail that *Simone is not in the kitchen*, and I am quite capable of grasping that entailment. So, it is *not* e-possible for me that *Simone is in the kitchen*, because, even though I don't know that she is not in the kitchen, some of the things I do know at the time, viz. (a) and (b), *obviously* entail that she is not in the kitchen.
- ¹⁵ Where: $\Diamond$ = It is possible that ...; K = I know that ...; b = I am a handless BIV; h = I have hands.
- ¹⁶ Where: $\Diamond_\mu$ = It is metaphysically possible that ...
- ¹⁷ To think otherwise is the epistemic equivalent of the naturalistic fallacy.
- ¹⁸ I am not begging the question against the skeptic here. I am simply making a conceptual point that from the mere fact that Bob knows that it is μ -possible that Jim is in New Orleans, it does not follow that Bob does not know that Jim is not in New Orleans, for as we have seen before, Bob may know that the μ -possibility in question is not actual.
- ¹⁹ Where: $\Diamond_e$ = It is e-possible that ...
- ²⁰ Many students, of course, when first presented with an argument like EAP2, feel compelled to accept the conclusion, because they cannot find anything wrong with the argument.

- ²¹ It is not idiosyncratic of D_3'' that it entails that S 's not knowing that $\sim p$ is a necessary condition for p 's being e-possible for S . In fact, virtually every purported account of e-possibility in the literature implies that S 's not knowing that $\sim p$ is necessary for p 's being e-possible for S .
- ²² DeRose defends a similar principle in DeRose (1991, pp. 596–601).
- ²³ Infallibilistic intuitions clearly underlie Saul Kripke's *a priori* Cartesian certainty account of e-possibility. See Kripke (1980, p. 143, fn. 72), where he proposes the following: p is e-possible for S iff S 's evidence does not justify *a priori* Cartesian certainty that $\sim p$.
- ²⁴ Knowledge_i and knowledge_f are distinguished as follows: Knowledge_i requires infallible justification, i.e. justification that entails that which it justifies. S is justified_i in believing that p [$J_i p$] only if p . Knowledge_f can be analyzed as follows:

$$(K_i) \quad K_i p \equiv (p \ \& \ Bp \ \& \ J_i p).$$

According to fallibilism, the kind of justification needed to convert true belief to knowledge_f must only make probable, but need not entail, that for which it is justification. As a result, fallibilism entails: $\Diamond(J_f p \ \& \ \sim p)$. This possibility and the closure principle with respect to justification together entail numerous Gettier-possibilities, including:

$$\Diamond[Bp \ \& \ J_f p \ \& \ J_f(p \rightarrow q) \ \& \ B(p \rightarrow q) \ \& \ J_f q \ \& \ Bq \ \& \ q \ \& \ \sim p \ \& \ \sim K_f q]$$

The latter possibility obtains when, as Gettier illustrated, S has a justified_f true belief that q which falls short of knowledge_f because S 's justification_f for q [to wit, $Bp \ \& \ J_f p \ \& \ J_f(p \rightarrow q) \ \& \ B(p \rightarrow q)$] fails to be appropriately connected to q 's truth and thus is defective (Gettier, 1963, pp. 121–123). Since fallibilism entails these possibilities, a fourth condition must be added to the traditional analysis of knowledge to rule out Gettier cases as instances of knowledge. For our purposes, the following condition will suffice: S is not Gettierized with respect to p [$\sim Gp$]. Accordingly, we can analyze knowledge_f as follows:

$$(K_f) \quad K_f p \equiv (p \ \& \ Bp \ \& \ J_f p \ \& \ \sim Gp).$$

- ²⁵ D_3'' should be revised accordingly:

D_{3f}'' p is e-possible_f for S at t iff (i) S does not know_f that $\sim p$ at t ; (ii) if one or more of the propositions S knows_f at t entail that $\sim p$, then it is not within S 's cognitive capacity at t to grasp that entailment; and (iii) if $\sim p$ is true and if one or more of the propositions S knows_f at t provide an adequate justificatory_f basis for believing that $\sim p$, then it is not within S 's cognitive capacity at t to see that these propositions justify_f her in believing that $\sim p$.

D_{3i}'' p is e-possible_i for S at t iff (i) S does not know_i that $\sim p$ at t , and (ii) if one or more of the propositions S knows_i at t entail that $\sim p$, then it is not within S 's cognitive capacity at t to grasp that entailment.

An infallibilistic version of condition (iii) of D_3'' is redundant where e-possibility_i is concerned.

- ²⁶ Where: $\Diamond_{ei}$ = It is e-possible_i that ...; K_i = I know_i that ...
²⁷ Where: $\Diamond_{ef}$ = It is e-possible_f that ...; K_f = I know_f that ...
²⁸ It is precisely such reasoning that inclines us to accept premise 1 of the original ambiguous EAP2.
²⁹ We might make the point as follows. Because EAP2f is valid, so is the following argument:
- $$\begin{array}{l} 2. \Diamond_{ef}b \rightarrow \Diamond_{ef} \sim h \\ 3. \Diamond_{ef} \sim h \rightarrow \sim K_f h \\ \sim 4. K_f h \text{ (i.e. } \sim \sim K_f h) \\ \therefore \sim 1. \sim \Diamond_{ef}b \end{array}$$

The above argument demonstrates that the falsity of EAP2f's conclusion entails the falsity of EAP2f's first premise, since, as we have seen, EAP2f 2 and 3 are true.

- ³⁰ Peter Klein makes a similar point with respect to skeptical arguments predicated on the closure principle. He claims that such arguments "virtually beg the question" because one of the premises in closure-based skeptical arguments can only be supported by a subargument that employs the conclusion of the main skeptical argument as a premise. See Klein (1995). While Klein properly diagnoses one way skeptical arguments can go wrong, he does not explain the source of their intuitive appeal, nor does he acknowledge the role equivocation plays in motivating skepticism.

REFERENCES

- Brendel, E.: 2003, 'Was Kontextualisten nicht wissen', *Deutsche Zeitschrift für Philosophie* **51**(6), 1015–1032.
Cohen, S.: 1988, 'How to Be a Fallibilist', *Philosophical Perspectives* 2 (Issue Epistemology), 91–123.
DeRose, K.: 1991, 'Epistemic Possibilities', *The Philosophical Review* **100**(4), 581–605.
DeRose, K.: 1995, 'Solving the Skeptical Problem', *The Philosophical Review* **104**(1), 17–52.
DeRose, K.: 1999, 'Contextualism: An Explanation and Defense', in J. Greco and E. Sosa (eds), *The Blackwell Guide to Epistemology*, Blackwell, Malden, MA, Oxford, 187–205.
DeRose, K.: 2000, 'Now You Know It, Now You Don't', *Proceedings of the Twentieth World Congress of Philosophy* Vol. 5, Epistemology, Philosophy Documentation Center, Bowling Green, Ohio, 91–106.
Gettier, E.: 1963, 'Is Justified True Belief Knowledge?', *Analysis* **23**, 121–123.
Hacking, I.: 1975, 'All Kinds of Possibility', *The Philosophical Review* **84**, 321–337.
Hume, D.: 1739, *A Treatise of Human Nature*, Second Selby-Bigge's edition (1978), Clarendon, Oxford.
Klein, P.: 1995, 'Skepticism and Closure: Why the Evil Genius Argument Fails', *Philosophical Topics* **23**, 213–236.
Kripke, S.: 1980, *Naming and Necessity*, Harvard University Press, Cambridge, MA.

- Lehrer, K.: 1971, 'Why Not Scepticism?', *The Philosophical Forum* 2.3, pp. 283–298 [reprinted in G. Pappas and M. Swain (eds), *Essays on Knowledge and Justification*, Cornell University Press, Ithaca and London, pp. 346–363, page references to the Pappas and Swain volume.]
- 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(4), 549–567.
- Nozick, R.: 1981, *Philosophical Explanations*, Harvard University Press, Cambridge, MA.
- Stroud, B.: 1984, *The Significance of Philosophical Skepticism*, Clarendon Press, Oxford.
- Vogel, J.: 2005, 'The Refutation of Skepticism', in M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology*, Blackwell, Malden, MA, Oxford, 72–84.

Department of Philosophy
Northern Illinois University
De Kalb, IL 60115
USA
E-mail: mylan-engel@niu.edu

GILBERT SCHARIFI

CONTEXTUALISM AND THE SKEPTIC:
COMMENTS ON ENGEL

ABSTRACT. Mylan Engel's paper (2004) is divided into two parts: a negative part, criticizing the 'costs of contextualism' and a constructive part proposing a 'non-contextualist resolution of the skeptical problem.' I will only address the constructive part here. The constructive part is composed of three elements: (i) a 'reconstruction' or 'reformulation' of the original skeptical argument, which draws on the notion of *epistemic possibility* (e-possibility), (ii) a distinction between two senses of 'knowledge' (and two corresponding kinds of e-possibility): *fallibilistic* and *infallibilistic*, and (iii) an argument which tries to hoist the skeptic by their own petard, namely the closure principle (CP). As I will argue, there are two ways to understand Engel's anti-skeptical argument. Only in one interpretation does the argument depend on the proposed 'reconstruction' of the skeptical argument in terms of e-possibility. But this version of the argument is unsound. More importantly, the skeptic has a strong *prima facie* objection at her disposal, which applies to both interpretations of the argument. If this objection is valid, Engel's argument does not hold. But once it is invalidated, his argument is superfluous.

1. THE SKEPTICAL PROBLEM

The 'skeptical problem' is created by the existence of a certain type of argument – the so called 'skeptical argument.' Engel gives the following example (BIV):

$$\begin{aligned}\sim K \sim b \\ \sim K \sim b \rightarrow \sim Kh \\ \therefore \sim Kh^1\end{aligned}$$

(Key: 'K' = 'I know that ...,' 'b' = 'I am a handless brain in a vat,' 'h' = 'I have hands.')

The skeptical problem consists in the fact that in arguments such as the above, the premises are apparently true, while the conclusion is apparently false.

2. SKETCH OF ENGEL'S RESOLUTION OF THE SKEPTICAL PROBLEM

Engel approaches this problem in two steps.

(I) In step one, Engel claims that the premises as well as the conclusion of the skeptical argument are equivocal. As he tells us, we



have to distinguish two senses of 'know': a strong sense, knowing *infallibly* ('know_i'), and a weak sense, knowing *fallibly* ('know_f'). If all instances of 'know' in the argument are read as 'know_i' the skeptic is, according to Engel, perfectly *right*: I do not know_i that I am not a brain in a vat (biv). I do *not* know_i that I have hands (i.e., the conclusion of BIV_i is, in fact, true). But, so we are told, we should not mind that we do not *infallibly* know such propositions.

Rather the important point is that given the fallibilistic reading, the skeptic, Engel claims, is entirely *wrong*: I do know_f that I have hands. And: I do know_f that I am not a biv! So the first premise of BIV_f is claimed to be false. In the second step, Engel attempts to provide an argument for this latter claim, which I will refer to as the 'anti-skeptical argument' (see below).

The first step of Engel's 'noncontextualist' alternative sounds a lot like contextualism. So what we get is a variant or a borderline case of contextualism, rather than an alternative to it. There are, of course, differences to (standard-) contextualism: According to Engel, 'know' is an *ambiguous expression* with just two senses. Engel explains the initial force of the skeptical reasoning by our not noticing this *ambiguity*. Whereas according to (standard-) contextualism, 'know' is a *context-sensitive term*, allowing a wide variety of different standards due to different contexts. (Standard-) contextualism explains the initial force of skepticism by our not noticing *switching standards* or *contexts*.

But these differences are not improvements: Contextualism can – at least presumably – handle a wide range of cases, e.g., Dretske's *zebra scenario* or the so called '*bank-cases*'. By contrast, Engel's proposal is of no help here, because these cases involve different standards for knowledge ascriptions, which are all below the infallibility level. (Engel could of course introduce additional senses of 'know' in order to account for these cases, too. But, if we are forced to choose between an ambiguous expression with a lot of (maybe countless) senses and a context-sensitive term, apparently the latter option is more reasonable.)

So far, Engel's account does not differ much from contextualism. But, whereas Engel and the contextualist both take the skeptic's claims to be wrong, insofar as the weak sense of (resp. a suitable low standard for) 'know' is concerned, they differ in their reasons for doing so. Here, Engel offers a genuine argument against the skeptic.

(II) When turning to Engel's anti-skeptical argument, things become more complicated: Engel's argument does not directly address BIV_f but a 'reformulation' or 'reconstruction' of it, which he calls 'EAP2f.'

Rather than in terms of knowledge, EAP2f is formulated by use of an epistemic operator ' $\Diamond_e$ ', which has to be read as: 'it is an epistemic possibility for me that ...' (For the moment, we have to rely on an intuitive understanding of this notion. A definition will be given below.) In accordance with the $\text{know}_i/\text{know}_f$ distinction, one has to distinguish $\Diamond_{ei}$ and $\Diamond_{ef}$. So EAP2f reads as follows:

$$\begin{aligned} &\Diamond_{ef} b \\ &\Diamond_{ef} b \rightarrow \Diamond_{ef} \sim h \\ &\Diamond_{ef} \sim h \rightarrow \sim K_f h \\ &\therefore \sim K_f h \end{aligned}$$

(Since in what follows, both 'know' and 'e-possibility' are always meant in the fallibilistic sense if not further specified, I will drop the indices).

The rationale behind this 'reformulation' seems to be the assumption that we are inclined to accept $\sim K \sim b$ (i.e., the first premise of the BIV-argument) *because* (!) we tend to accept two premises from which it follows (i.e.: $\Diamond_e b$ and $\Diamond_e b \rightarrow \sim K \sim b$) (cf. Engel, 2004, p. 214). If this assumption were true, EAP2 would count as a piece of analysis, an 'elaboration' of BIV.

This becomes more transparent if we rewrite EAP2 as follows:

$$\begin{aligned} &\Diamond_e b \\ &\Diamond_e b \rightarrow \sim K \sim b \\ &\sim K \sim b \rightarrow \sim K h \\ &\therefore \sim K h^2 \end{aligned}$$

Engel presents his anti-skeptical argument as follows:

[T]he skeptic is in no position to assert that its first premise is true, for *suppose* that I know_f that I have hands (h). Then I know_f a proposition – namely, h – that entails that I am not a handless BIV. ... So, the skeptic can only rationally assert EAP2f's first premise ... if she *assumes* the truth of EAP2f's conclusion". (Ibid., p. 221)

As I will argue, there are two possible interpretations of Engel's anti-skeptical argument, which I call the 'straightforward' interpretation and the 'sophisticated' interpretation. Only in the second interpretation does the argument in fact depend on the proposed 'reconstruction' of BIV.

3. THE STRAIGHTFORWARD INTERPRETATION

In this interpretation, the essence of the argument is captured by the following dialogue:

Skeptic: You don't have any evidence for not being a biv, so I am entitled to hold the first premise ...

Me: Stop. I have evidence: (I know that) I have hands. That I have hands entails that I am not a biv. (So, by CP, I know that I am not a biv).

Skeptic: That does not work, since you don't know that you have hands.

Me: Ha, got ya: You just assumed the conclusion.

So the skeptic would have to settle the question of whether I do or do not know that *h* before she can state her first premise. This would be putting the cart before the horse.

Note that this argument in no way depends on the 'reconstruction' of the original skeptical argument in terms of *e*-possibility instead of knowledge. (It does not make any difference to this argument whether the skeptic's first premise reads 'It's *e*-possible for me that *b*' or 'I do not know that $\sim b$ '). The several pages Engel spent on the introduction of the notion of *e*-possibility were unnecessary and could be left out altogether.

So in the next section, I will attempt to see how this argument could be understood in a way that ascribes the notion of *e*-possibility an essential role.

4. THE SOPHISTICATED INTERPRETATION

The first premise in the skeptic's argument EAP2 reads 'It is *e*-possible for me that *b*.' So what exactly does the skeptic claim in asserting the first premise? The definition by which the notion of *e*-possibility is introduced says:

p is *e*-possible for *S* at *t* iff: (i) *S* does not know that $\sim p$, and (ii) *S* could not come to know that $\sim p$ at *t*, strictly on the basis of propositions *S* knows at *t*. (cf. *ibid.* p. 220)

This definition is elaborated on in a footnote on the same page. For the purpose at hand, I will simplify the elaborated definition as follows:

p is *e*-possible for *S* iff

(ep) it is not the case that the propositions known by *S* entail $\sim p$ or justify the belief that $\sim p$.³

So in asserting the first premise, the skeptic claims that the relevant instance of (ep) is true: I do not know any proposition that entails the

proposition $\sim b$ or justifies the corresponding belief. This *a fortiori* means that I do not know the proposition h . So in order to establish the first premise, the skeptic has to establish *inter alia* that I do not know that h .

Now consider Engel's claim that the skeptic cannot "give an independent reason for thinking EAP2f's first premise true [i.e. a reason that does not make reference to the truth of EAP2f's conclusion] ... because *e-possibility_f is analyzed in terms of knowledge_f*" (ibid., p. 222, my italics).

Following the sophisticated interpretation, Engel claims here that the skeptic cannot give an independent reason, because *knowledge* is the prior notion and the notion of *e-possibility* is defined and introduced in terms of it.⁴

Note: If we contrapose the skeptic's argument, we do not face the same problem:

I know h .

$\therefore$ It's not e-possible for me that b .

Here the premise is stated in terms of knowledge. Knowledge does not have to be analyzed in terms of e-possibility, because it is the prior notion.

The 'sophisticated interpretation' poses the following problem: There is obviously a tension between (i) claiming that there can be no independent reason for the first premise, because e-possibility has to be analyzed in terms of the *prior* notion of knowledge, and simultaneously claiming (ii) that EAP2 is a warranted reconstruction of the original skeptical argument. If we accept $\sim K \sim b$, *because* we accept $\Diamond_e b$ and $\Diamond_e b \rightarrow \sim K \sim b$, this makes *e-possibility* the 'prior' notion.

I will not expand on this point because there are objections that apply to both versions of the argument.

5. THE STANDOFF-PROBLEM

At best, Engel gets a *standoff*: Neither can the skeptic legitimately claim that we do not know such ordinary propositions as h , nor can we claim that we do. You might think of this situation as follows: The skeptic has launched an assault to the legitimacy of our knowledge claims. This could be warded off. So everything is just fine. But this would be a misdescription of the situation. In fact we are left with a threatening, true conditional of which we do not know how to

evaluate the antecedent: If we do not know $\sim b$, then we do not (cannot) know such ordinary propositions as h . So we just do not know whether we (can) know such propositions or not. This is surely not enough to “account for the truth of our commonsense judgement that we do in fact possess lots of ordinary knowledge” – as Engel promises to do (ibid., p. 203).

6. THE INDEPENDENT-REASON-OBJECTION

Contrary to what Engel claims, there is indeed ‘an independent reason’ the skeptic can cite in favour of his first premise, *even* if we grant that e-possibility must be analyzed in terms of knowledge. Strongly supported by intuition, the skeptic may claim that sensitive belief is a necessary condition for knowledge. Call this the sensitivity requirement (SR). Roughly: A belief that p is *sensitive* only if the subjunctive conditional ‘*If p were not true, I would not believe that p* ’ is true.⁵

So to dispute the ‘straightforward version’ of the argument the skeptic may claim that I do not know that $\sim b$, because the belief that $\sim b$ is not sensitive (this simply is the way all radical skeptical hypotheses are constructed).

Turning to the ‘sophisticated version’ of the argument: even if we grant that EAP2 is a correct reconstruction of the original skeptical argument and also grant that e-possibility has to be analyzed in terms of knowledge, this does not prevent the skeptic from establishing his first premise, e.g. that it is e-possible for me that b .

Remember the definition of e-possibility:

p is e-possible for S at t iff: (i) S does not know that $\sim p$, and (ii) S could not come to know that $\sim p$ at t , strictly on the basis of propositions S knows at t .

The relevant instance of (i), i.e. I do not know that $\sim b$, is already shown to be true by invoking SR. In order to show that b is e-possible for me, the skeptic now has to show that the relevant instance of (ii) is also true. He can do so by a short indirect argument:

Suppose that this instance of (ii) is not true.

That means: I *could* come to know that $\sim b$.

Thus: It’s (metaphysically) possible that I know that $\sim b$.

Thus: It’s (metaphysically) possible that I have the sensitive belief that $\sim b$.

But that is impossible, the skeptic will argue, so we have a *reductio*. Sure: I could believe that $\sim b$, but such a belief could never be sensitive.⁶

7. THE SENSITIVITY REQUIREMENT AND THE CLOSURE PRINCIPLE

Peter Klein (2002) objects that the skeptic cannot successfully appeal to SR. Klein argues as follows: As generally acknowledged, CP fails if SR holds. So by invoking SR, the skeptic gains a reason for his first premise, but loses the warrant for his middle premise (i.e. $\sim K \sim b \rightarrow \sim Kh$).

Why is CP supposed to fail if SR holds? Take c to be any contingent proposition (*cogito* cases aside). From ‘I know that c ’ it follows by CP that I know the denial of a suitably chosen skeptical hypothesis (‘I am in a $\sim c$ skeptical scenario’). By SR it follows that I have the *sensitive* belief that I am not in a $\sim c$ skeptical scenario. But that is impossible. So, we are told, the assumption that CP and SR are *both* valid cannot be right.

But what really follows is that it cannot be true that CP and SR are both valid *and* I know that c . But that no one can know that c is just the claim the skeptic wants to defend.⁷ In the dispute with the skeptic this argument is a *petitio*. You hardly can argue against the skeptic by saying: Well, CP and SR are heavily supported by intuition, but you are only allowed to use one of them, because together they would deliver the result you are aiming at. That would be rejecting a *counterargument* to your position, on the sole ground that it is a counterargument to your position.

(Note that SR only gives a necessary condition for knowledge. Acceptance of SR does not amount to the acceptance of the claim that (true) sensitive belief is sufficient for knowledge, which is incompatible with CP.)

In discussion, an objection was made that we do not have a *petitio* here, because it is not claimed that I *in fact* know that c , but only *supposed* that I do so. Now, *suppose* I know that there is a male vixen. If knowledge is factive, it follows that there is a male vixen. But this is impossible, so knowledge is not factive. This clearly is not a good argument, because it supposes something impossible. It would be a good argument if one could show that what it supposes is in fact possible. Correspondingly, to make the above argument against the skeptic a good one, one first would have to show that it is really possible that I know c – which is exactly the point at issue. It does not matter here whether the ‘impossibility’ is located in the content of the purported knowledge or in the purported fact that some content is known. All that matters is that one can infer by valid principles the impossible from the impossible.

8. THE CONTEXTUALIST SOLUTION

The contextualist's answer to the skeptic has to be understood as a two-step argument. Roughly, mixing elements from David Lewis (1996) and Keith DeRose (1995), it amounts to this:

The first step is to introduce different standards for knowledge ascriptions. According to a *low standard*, I may count as knowing *h*. The standard is called 'low' because it allows for the *presupposition* of a set of relevant propositions *Q*, including the denial of the skeptical hypothesis, i.e. $\sim b$.⁸ So the claim that I know that *h*, roughly equals the claim that my evidence eliminates every $\sim h$ -possibility, except those which are 'properly ignored', which means ruled out *ex ante* by means of the presupposition.

Apparently, my *epistemic position* with respect to *Q* (or an element of it, including $\sim b$) is at least as strong as my epistemic position with respect to *h*. So, if I count according to a low standard as knowing that *h*, I should count by the very same standard as knowing that $\sim b$, since my evidence eliminates every *b*-possibility except those which are ruled out *ex ante* by means of the presupposition $\sim b$. This, of course, does not leave any possibilities for my evidence to eliminate.

The last claim clearly sounds strange. It may be tenable to claim 'I know that *h* according to a low standard of knowledge which allows for the presupposition of $\sim b$ '. But 'I know that $\sim b$ according to a low standard of knowledge which allows for the presupposition of $\sim b$ ' is odd. The contextualist has to maintain that despite its oddness, the second claim is not false. So, he has to provide a non-threatening explanation for its apparent oddness. This is the second step of the argument. DeRose assumes that such a claim sounds odd, because we *falsely* think that sensitive belief is a necessary condition for knowledge. He proposes a rival principle to explain the intuitions which normally mislead us to hold SR true. He calls this principle the *rule of sensitivity*. While SR belongs to semantics, the rule of sensitivity is taken to be a pragmatic principle which governs the contextually relevant standards of knowledge ascriptions. Roughly: Whenever a knowledge ascription '*S* knows that *p*' is at issue, by the rule of sensitivity, the standards of knowledge tend to rise up to a point where only a sensitive belief would count as knowledge. So, the rule of sensitivity is supposed to back up the claim that SR is not a necessary condition for knowledge in spite of the fact that we always tend to judge knowledge ascriptions *as if* it were.

To locate the points of disagreement between the contextualist and the skeptic, it may be useful to look at how the skeptic might respond to the contextualist.

The skeptic may resist the first step of the argument, maintaining that ‘know’ is an *absolute term*: there is only one (ultimate high) standard associated with it. (He may concede, in an Unger-like fashion, that contextual standards in fact govern the appropriateness of a knowledge ascription, while denying that they are relevant to its truth-conditions).

But the skeptic may as well, if only for the sake of the argument, allow different standards for ‘know’. Insisting on SR (and CP) he can still argue that no one knows a proposition like *h*: Regardless of which standard of ‘know’ is employed, no one knows $\sim b$, since no one can have *any* evidence for this belief; it is necessarily *insensitive*. So no one knows *h* either.

9. THE CONTEXTUALIST’S AND ENGEL’S SOLUTION

We have two theses about the correct semantics of ‘know,’ which mark the disagreement between the skeptic and the contextualist: (1) ‘Know’ is an absolute term. (2) Sensitive belief is a necessary condition for knowledge.

The skeptic must hold at least one of them. Engel and the contextualist both deny (1): The contextualist by claiming that (1*) ‘know’ is a context-sensitive term; Engel by claiming that (1**) ‘know’ is ambiguous between a fallibilistic and an infallibilistic reading. Both take the skeptic’s claims (that I do not know . . .) to be wrong insofar as the weak sense of ‘know’ – respectively a suitable low standard for ‘know’ – is concerned. (And both claim that the skeptic is right, if done otherwise). As I argued, there is a good sense in which – despite the differences – Engel’s (1**) may be regarded as a variant of the contextualist’s (1*).

The relevant (contextualist’s) low standard is a standard which allows for the presupposition that skeptical hypotheses do not hold – call this ‘standard *L*’. If we assume (1), we arrive at the anti-skeptical solution ‘by definition’ – once (2) is dismissed. There will be no need for any further argument. The contextualist (typically) challenges (2) by appealing to a rival principle, namely the rule of sensitivity. But this move is not essential to the contextualist’s rebuttal of the skeptic. Any other successful argument against (2) would do, too.⁹

As I argued, (2), i.e. SR, provides a strong *prima facie* objection against Engel's anti-skeptical argument. Engel may of course reject (2) despite its intuitive force. (He may do so by appealing to the rule of sensitivity or by any other argument). But, as it seems to me, once (2) is dismissed, there will be no longer any need for Engel's anti-skeptical argument. If (2) is dismissed, (1**) provides the aimed at anti-skeptical solution without any further ado. If (1*) provides the contextualist with all he needs (once (2) is dismissed), (1**) should do so for Engel. To sum it up: If SR holds, Engel's argument does not hold. If SR can be dismissed, Engel's argument is superfluous.

Here one might object that in contrast to the contextualist, Engel denies (1) in a way that does not render his argument futile. Engel's distinction between two senses of 'know' draws on the notion of *fallibility*. One might argue that although the notion of *fallible knowledge* is somehow associated with some 'low' standard, it is not associated with (the contextualist's) standard *L*. Fallible knowledge, so one might argue, invokes a standard *above* or just *different* from standard *L*; at any rate, a standard which does not presuppose that the skeptical hypothesis do not hold. (After all, the distinctive characteristic of fallible knowledge is that it only requires justification which makes probable that for which it is a justification.) From this perspective, the anti-skeptical solution is not arrived at 'by definition'; we need a further argument.

To substantiate this objection, however, one would have to show that there is some reasonable understanding of *fallible knowledge* that does not invoke a standard which presupposes the non-actuality of skeptical scenarios. That means, one has to show that a reasonable understanding of fallible knowledge does not connect this notion with standard *L*. I doubt that this is possible. In the most demanding sense, fallible knowledge that *p*, would require the best possible empirical evidence for *p*. But, strange cases aside, in a skeptical scenario, there is hardly any empirical evidence for anything.

But even if one could show that fallible knowledge in fact differs in the alleged way from knowledge according to the contextualist's low standard *L*, there is no need for Engel's anti-skeptical argument: Once he has denied (1) by invoking a distinction between fallible and infallible knowledge, nothing should prevent a further distinction introducing 'knowledge according standard *L*,' since this further distinction would be at least as motivated and well-founded as the original one.

ACKNOWLEDGMENTS

I would like to thank Thomas Grundmann for helpful comments on an earlier draft of this paper, and Mylan Engel for help with the language.

NOTES

- ¹ Here, as well as in the following, it is assumed that the person in question knows that the relevant entailment holds (here: that $\sim b$ follows from h); otherwise the second premise would not be warranted by CP.
- ² Engel should not have any quarrels here, since he accepts CP.
- ³ What Engel really has in mind is rather the following: It is not the case that both the propositions known by S entail $\sim p$ or justify the belief that $\sim p$. And it is within S 's cognitive capacity to grasp that entailment/justification relation. Since the second conjunct does not matter in what follows, I dropped it to simplify matters. For the same reason, I have omitted any temporal references.
- ⁴ I understood Engel to say that there can be no independent reason for the skeptic's premise because: "... (fallibilistic) *e-possibility* is analyzed in terms of (fallibilistic) *knowledge*." In this case, the second interpretation of the argument is the one Engel really has in mind. But one can also read the quote as claiming that there can be no independent reason because "... fallibilistic *e-possibility* is analyzed in terms of fallibilistic knowledge."
- ⁵ A more qualified explanation of *sensitive belief* is proposed by DeRose (1995), roughly: In order to be sensitive, a belief as to whether p must be 'truth tracking' – i.e. match the fact of the matter as to whether p is true – over a range of possible worlds including at least the closest $\sim p$ worlds.
- ⁶ Note: the negation of (ii) has to be understood roughly as 'would S think for a minute, she would *actually know* that $\sim p$.' So, for the relevant instance of non-(ii) to be true there not only has to be a possible world in which I have the sensitive belief $\sim b$, but a very close world in which this happens. To contradict this, the skeptic only has to defend the more modest claim, that *among the very close worlds*, there is none in which I have such a belief.
- ⁷ The skeptic may nevertheless hold a version of SR which allows for knowledge of analytic propositions, since the subjunctive conditional '*If p were not true, I would not believe that p* ' is true by virtue of a vacuous antecedent for any necessarily true p .
- ⁸ Here I use 'presupposition' in a rather loose sense.
- ⁹ As long as it is not incompatible with (1*), of course.

REFERENCES

- DeRose, K.: 1995, 'Solving the Skeptical Problem', *Philosophical Review* **104**, 1–52.
 DeRose, K.: 2000, 'Now You Know It, Now You Don't', *Proceedings of the Twentieth World Congress of Philosophy*, Vol. V, 91–106.
 Engel, M.: 2004, 'What's Wrong with Contextualism', *Erkenntnis*, **61**, 203–231.
 Klein, P.: 2002, 'Skepticism', in P. K. Moser (ed.), *The Oxford Handbook of Epistemology*, Oxford.

Lewis, D.: 1996, 'Elusive Knowledge', *Australasian Journal of Philosophy* **74**, 549–567.

Philosophische Fakultät I
Philosophisches Institut
Universität des Saarlandes
Postfach 151150
66041 Saarbrücken
Germany
E-mail: scharifi@gmx.de

BRUCE RUSSELL

HOW TO BE AN ANTI-SKEPTIC AND A NONCONTEXTUALIST

ABSTRACT. Contextualists often argue from examples where it seems true to say in one context that a person knows something but not true to say that in another context where skeptical hypotheses have been introduced. The skeptical hypotheses can be moderate, simply mentioning what might be the case or raising questions about what a person is certain of, or radical, where scenarios about demon worlds, brains in vats, The Matrix, etc., are introduced. I argue that the introduction of these skeptical hypotheses leads people to fallaciously infer that it is no longer true to say that the relevant person knows. I believe that that is a better explanation of the so-called intuition that the person does not know than the contextualist's who claim that raising these skeptical hypotheses changes the standards that determine when it is true to say "S knows that P." At the end I raise the possibility that contextualists might defend their view on pragmatic rather than skeptical grounds by arguing that the standards of evidence rise when more is at stake in a practical sense.

Many arguments for contextualism proceed by way of example. The examples are supposed to generate intuitions about what people know and what they don't know. Contextualism is then defended on the grounds that it can best explain these intuitions. Keith DeRose thinks that examples taken from nonphilosophical conversation are at least as important to the defense of contextualism as those taken from philosophical discussions of skepticism. In the ordinary cases, moderate skeptical hypotheses will be introduced regarding what might be the case (say, that a bank will be closed on Saturday), what a person is, or is not, certain of (e.g., that the bank will be open), etc., but radical skeptical hypotheses involving brains in vats, evil demons, The Matrix, and the like do not play a role.¹ Below is an example from a nonphilosophical situation that introduces a moderate skeptical hypothesis and one that introduces a radical skeptical hypothesis. The first is based on an example of DeRose.

Hannah and her husband are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoon. Thinking that it isn't very important that their paychecks are deposited right away, Hannah says "I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit them tomorrow



morning.” But then Hannah’s husband reminds her that a very important bill comes due on Monday, and that they have to have enough money in their account to cover it. “Banks do change their hours. Are you certain that’s not what is going to happen tomorrow?” Hannah concedes, uttering “I guess I do not really know that the bank will be open tomorrow.”

And here is another example.

Prof *X* asks her student whether she knows that she is in a classroom to which she replies, “Of course, I know!” Prof *X* responds by asking the student whether she knows that she is not a brain in a vat. After thinking about it for a moment, the student concludes, “I guess I don’t *really* know I’m in a classroom.”²

In the bank case the intuitions are that Hannah knows before her husband asks her if she is certain the bank will be open on Saturday and then does not know once he has asked the question.³ In the classroom case the intuitions are that the student knows she is in the classroom before her professor asks whether she knows she is not a brain in a vat but does not know once he has asked that. The contextualist’s explanation is that some contexts are low-standard contexts and others are high-standard ones. When one shifts to a high-standard context, one may no longer have sufficient evidence to qualify as knowing some proposition even though that level of evidence would have sufficed for knowledge in a low-standard context.

There are two different ways to interpret these examples. On one interpretation, once the question is raised doubt ensues. On this interpretation, once her husband raises the question Hannah doubts whether the bank will be open on Saturday and the student doubts that she is in a classroom once her professor has raised the possibility that she does not know that she is not a brain in a vat. Of course, on this interpretation neither Hannah nor the student have knowledge after the relevant question is raised because knowledge requires belief and neither of them believe the proposition they had believed earlier.

However, in DeRose’s original bank example it is stipulated that the subject in the example, namely, DeRose himself, remains confident that the bank will be open on Saturday. So it is assumed that the subject does not come to doubt that the bank will be open on Saturday and still believes that it will be open. When DeRose’s wife asks him whether he knows that the bank will be open the next day, he responds, “Well, no. I’d better go in and make sure.” (DeRose, 1992, p. 913). Here DeRose is like the person about to go on vacation who gets out of the car and goes back into the house to make sure the stove has been turned off and the doors locked even though he

checked a couple of times before heading for the car. The person still believes the stove is off and the doors have been locked even though he still wants to make sure.

If we have the “intuition” that DeRose does not know the bank will be open on Saturday, I believe it is because we implicitly infer: (a) if he knows it will be open on Saturday, then it must be open then; (b) but it might not be open; so (c) he doesn’t know it will be open. We see his “making sure” as ruling out (b). Once that is done, then this argument won’t go through, and he will once again know that the bank will be open. Or as the contextualists prefer to say, the sentence “DeRose knows the bank will be open on Saturday” is true once he goes back in and checks, just as it was before his wife raised her question.

Of course, (a) is a misinterpretation of “What you know must be true” applied to the case of the bank’s being open. The correct interpretation applied to that case is: (a*) necessarily, if he knows it will be open on Saturday, then it will be open then, not (a) which has the necessity in the consequent of the conditional. Because it is easy to confuse (a) with (a*), we are epistemically blameless in believing (a) and, in that sense, justified in believing (a).⁴ So there is a sense in which we are justified in believing the conclusion, namely, that DeRose does not know the bank will be open. Of course, if our belief that DeRose does not know the bank will be open on Saturday is based on this implicit inference, then we do not really have *an intuition* that he doesn’t know this since an intuition is non-inferential. Still, it could explain how we are justified in believing DeRose doesn’t know that the bank will be open on Saturday.

Granted that before scrutinizing the relevant inference we are justified in believing that he doesn’t know. Still, is it true that he doesn’t know? I believe it is not; I think DeRose knows the bank will be open on Saturday if he knew that before his wife raised the question.

Imagine that what is at stake is not, say, making sure there is enough money in the bank on Monday to pay the mortgage but making sure there is enough money there to pay a ransom to keep some kidnappers from killing DeRose’s son. Suppose DeRose looks at the bank hours posted on the door, and it says it will be open on Saturday. Doesn’t he then know that it will be open? I think most people’s intuition will be that he does know.

But suppose he returns to the car and his wife says, “Yes, but they might not have changed the posted hours when they changed the hours. Better go in and ask the manager just to make sure.” Here,

again, the implicit inference that makes us think he does not know it will be open is: “what he knows must be true, but the bank might not be open tomorrow, so he doesn’t know it will be.” This argument, of course, equivocates on “what he knows must be true,” and so is invalid. Still, we can be justified in believing the conclusion on the basis of the argument.

The best explanation of our judgment that DeRose does not know the bank will be open even after seeing the posted hours is that we are taken in by a fallacious argument, that is, an argument that looks good even though it is bad. He really does know it will be open even though we are justified in thinking he does not. Insofar as that is the best explanation in the case where he looks at the posted hours, a similar explanation is the best one in the original version of the example, too. Before his wife mentioned the possibility that the bank has changed its hours, he knew that it would be open on Saturday. After she raised that possibility, he still knew it would be open but, on the basis of a fallacious inference, he, and we, are justified in believing that he did not know.

What about the student who concludes that she does not know that she is in a classroom when asked by her professor if she knows she is not a brain in a vat? Here my intuition is that provided she has not come to doubt that she is in a classroom, of course, she knows she is, even if she thinks she does not. Now those who have the opposite “intuition” I believe make the following inference: (1S) the student does not know that she is not a brain in a vat, (2S) if she knows she is in a classroom, then she knows she is not a brain in a vat, so (3S) she does not know she is in a classroom.

Of course, they will be justified in accepting (3S) only if they are justified in accepting (1S) and (2S). They are justified in accepting (2S) because of closure: if you know *P* (that you are in a classroom) and know that *P* entails *Q* (that you are not a brain in a vat), then you know *Q*. Mylan Engel considers, and then criticizes, an Argument from Possibility (AP1), that might be used to justify (1S):

- (1) It is possible that I am a BIV (brain in a vat).
- (2) If it is possible that I am a BIV, then I don’t know that I am not a BIV.
- (3) Therefore, I don’t know that I am not a BIV.⁵

This argument can be thought of as a more exact formulation of the following more informal argument: (1*) for all I know, I’m a BIV and (2*) if for all I know I am a BIV, then I don’t know I am not a BIV, so (3*) I don’t know that I am not a BIV.

A case can be made that this informal argument equivocates on “for all I know.” In (1*), “for all I know” means “for all I know *for certain*, that is, from what is entailed by my evidence,” whereas in (2*) “for all I know” means “I do not know that I am not...,” in which case (2S) is a tautology. It becomes: if I do not know that I am not a BIV, then I don’t know that I am not a BIV. But the equivocation is subtle enough to justify a person who gives the argument in accepting the conclusion.

Engel’s discussion of a second similar Argument from Possibility (AP2) shows how the above argument involving possibility equivocates on both “possible” and “knows” (Engel 2004, pp. 218–222). Again, the equivocations are subtle enough to justify a person in accepting the conclusion that they do not know they are not a BIV.

So someone who has an “intuition” that the student does not know that she is in a classroom once the professor has raised the question of whether she knows she is not a brain in a vat will have a justified belief that the student does not know she is in a classroom only if he relies on some fallacious argument involving a subtle equivocation. Whether the so-called intuition is justified or not, it does not follow that the student really does not know she is in a classroom once the professor raises his question. As long as the student does not doubt that she is in a classroom, she still knows she is, even if she believes otherwise.⁶ The student should have responded, “Look, I might be in The Matrix, but I still know I am sitting in a chair. So while it’s true that if I know I am sitting in a chair, I know I am not in The Matrix, I can’t agree that I don’t know that I am not in The Matrix. Why think I don’t know that? So why think I don’t know that I am not a brain in a vat? Of course I know that I am not a brain in a vat. The best explanation of my experiences is that there are real chairs and people causing them, not some super scientists.”

When Engel analyzes an argument involving claims of possibility and knowledge like the one given above, he distinguishes two senses of “epistemically possible,” a fallible and an infallible one, and two senses of “knows,” a fallible and an infallible one. That means that there are four possible interpretations of the argument. He argues that one interpretation involves invalidity, another a false premise, and a third begs the question. The fourth is sound but concludes that we do not *infallibly* know that we are not BIVs, from the fact that our evidence does not entail that we are not.

This sort of skepticism is innocuous, and allowing that we lack this sort of knowledge is of no help to the contextualist. According to this sort of skepticism, in *no* context do we have infallible knowledge.

And as far as the argument goes, in *all* contexts fallible knowledge is possible, especially of ordinary things such as that I have hands and am in a classroom. Further, it allows that *in the same context* I can lack infallible, but possess fallible, knowledge of something, say, that I am in a classroom. Put metalinguistically, it allows that there is an interpretation of the sentence “I know I am in a classroom” that is false in some context and that there is an alternative interpretation of the very same sentence that is true in that same context. That is not possible for contextualism.

The trick for the person who opposes both skepticism and contextualism is to explain how the arguments that *seem* to support skepticism do not support any interesting version of it and yet are good enough to justify people in thinking that, say, Hannah and the student lack knowledge once the relevant questions are asked. Arguments for skepticism that contain subtle equivocations fit the bill: they provide justification for the skeptical conclusions without really providing sound arguments for any interesting form of skepticism.

I believe that nearly all arguments for skepticism fit this pattern, that is, they contain subtle equivocations. Gilbert Scharifi suggests that there is an argument for the claim that we do not know that we are not brains in vats that does not equivocate on “possibility.” Similar to Nozick, he holds that knowledge must, at least in some contexts, be “sensitive to the truth” in the sense that a person does not know something if she would still believe it if it were false (Scharifi, 2004, p. 238). So, for instance, the student in the philosophy class would not know that she is not a brain in a vat because she would still believe that she was not even if she were a brain in a vat. But as Alvin Plantinga has argued, if sensitivity to the truth were a requirement of knowledge, even if in only some contexts, it would be possible to know a conjunction yet fail to know one of its conjuncts.⁷ Let one of the conjuncts be the denial of some skeptical hypothesis, say, the denial of “I am a BIV.” “I am not a BIV” is not sensitive to the truth because I would still believe it even if I were a BIV. Let the other conjunct be some ordinary proposition, say, “I am seated in a chair (that is not in a vat!).” The conjunction could be sensitive to the truth because whenever it is false it would be because I am not seated in a chair, and I would not believe I that I was seated in a chair if I were not. And surely I could know, “I am not a BIV and I am seated in a chair,” given all the evidence I have and the fact that this conjunction is sensitive to the truth. Still, if sensitivity to the truth is a necessary condition for knowledge, at least in some contexts, then I

will not know that I am not a BIV; and thus, it would be possible for me to know a conjunction but not know one of its conjuncts. This is absurd.

So sensitivity to the truth is not a necessary condition of knowledge in any context. Perhaps sensitivity to the evidence is a necessary condition of knowledge. Justification is a necessary condition of knowledge, and even if I am in a demon world, or a brain in a vat, to be justified in believing, say, that some bird is a golden eagle I must be able to distinguish golden eagles from hawks. In other words, justification (and so knowledge) requires the ability to discriminate between evidence for *X* and evidence for *Y*, and in that sense requires sensitivity to the evidence. This ability to distinguish between what *looks like* an *X* and what *looks like* a *Y* is easily conflated with the ability to tell what is an *X* and what is a *Y*, that is, with sensitivity to the truth, because each could be described as the ability to tell the difference between *X*'s and *Y*'s. But while sensitivity to the evidence is required for justification, and hence for knowledge, sensitivity to the truth is not.

All of the arguments about not knowing that the animal you see in the zoo is a zebra because you cannot eliminate, or rule out, the possibility that it is a cleverly painted mule equivocate on "eliminate, or rule out." In normal circumstances the best explanation of what you see is that it is a zebra, and that allows you to "eliminate, or rule out" the possibility that it is a cleverly painted mule. Of course, its being the best explanation does not guarantee, does not entail, that the animal is not a cleverly painted mule. But why think that knowledge requires *that* sort of eliminating or ruling out?

Some arguments for skepticism are arguments for skepticism about knowledge of the future. I think some people know that they are going to be alive a week from today. They have enough evidence about their health and what their activities will be in the next week to know that they will be alive in a week. Suppose, however, someone offers them a life insurance policy of, say, one million dollars to cover that week alone, and they will sell it to that healthy person for ten dollars. It seems rational for them to spend the money, especially if they have young children and a spouse that depend on them. But, the argument goes, if they know they will be alive in a week, it would not be rational for them to pay even a dollar for the insurance. So they must not know that they will be alive in a week.

Perhaps this argument presupposes: if you know something is true, then it is certain that it is. And if it is certain that you will be alive in a week, it is irrational for you to pay anything for life insurance

covering only that week. So, if you know you will be alive in a week, then it is not rational for you to spend anything, even one dollar, for life insurance to cover only that week.

This argument equivocates on “is certain.” As was seen when discussing Hannah, knowledge requires lack of doubt, and in that sense certainty. Also, “if you know something is true, then it is certain that it is” might be a misleading way of saying, “What you know must be true.” But there is no reason to think that lack of doubt or the fact that what is known must be true requires that the statement of what is known be a necessary truth or entailed by the person’s evidence. However, the second premise in the argument requires this reading since if your being alive next week is entailed by your evidence (or is a necessary truth!), then it would be irrational to spend any money on life insurance. But otherwise it need not be irrational to purchase insurance.

“If you know you will be alive next week, it is not rational to pay anything for life insurance covering only next week,” seems true until you realize that knowledge does not require certainty and the lack of certainty alone can make it rational to purchase life insurance. As was just seen, an argument can be given for the conditional that links knowledge and irrationality, but it equivocates. So on reflection, this life insurance argument provides no reason for skepticism about the future. However, before spotting the false premise in that argument, or the equivocation in a sub-argument for that premise, the argument appears good enough to provide justification for the skeptical conclusion. But, as noted earlier, this provides no solace for either skeptics or contextualists. Bad arguments that appear good, that is, fallacious arguments for skepticism, neither support skepticism nor provide data that noncontextualists cannot explain.

Advice for anti-skeptical noncontextualists: in every case where a contextualist believes there is data that only his view can explain, find some bad skeptical argument that rests on either a moderate or radical skeptical hypothesis, where that argument appears good enough to support what the contextualist calls an intuition that knowledge is absent. Explain how it looks good and so provides justification for the intuition that knowledge is absent. Explain that my being justified in believing that *S* does not know that *P* does not entail that *S* does not know that *P*. Maintain that if the evidential situation of *S* remains unchanged and *S* believes that *P* across contexts, then if *S* knows that *P* in one context, he knows it in the other regardless of what the ascriber is justified in believing about *S*’s knowledge. Against the skeptic, explain how the argument the

ascriber relies on is really bad, despite its initial appearances. Point out the costs of contextualism, as Engel does. That is how to be an anti-skeptic and noncontextualist.

Advice to the contextualist: to avoid competition with my explanation, start with an example that does not introduce either a moderate or a radical skeptical hypothesis. Guard against fallacious skeptical inferences by giving several and explaining the fallacies committed. You might use the Hannah case varying the situation from the original where little is at stake, to the second scenario where much is at stake, to the third where a great deal is at stake. However, do not raise even moderately skeptical questions in the example for that would make my competing explanation viable. Second, argue for the view that justification is relative not only to the evidence a person has but to the evidence she *should* have.⁸ Third, argue that the evidence a person *should* have is relative to context, or perhaps to interest or perceived interest, not just to the nature of the proposition believed. Finally, join the noncontextualists in criticizing the arguments for skepticism. Contextualists, and their interest-theory cousins, should dissociate themselves from all forms of skepticism and rely solely on examples taken from nonphilosophical conversation that do not introduce any (not even moderate) skeptical hypotheses.

If my advice is followed, I do not know whether the noncontextualists or the contextualists will win. Put more generally, I do not know whether the following view (sometimes called invariantism) will turn out to be true: two different people, regardless of the context each is in, must judge another person as knowing, or not knowing some proposition, provided that their total evidence about that other person's situation is the same and they are in the same psychological state. Further, any given person must judge two different people as both having or both lacking knowledge if those two people base their beliefs on the same total evidence, are in the same psychological state, and believe the same proposition, even when making these two judgments in different contexts. For the invariantist, similar things will hold for justified belief. On invariantism, only differences in evidence or psychological state matter when it is a question of what ascribers are justified in believing (or know) about others, or a question of what those others are in fact justified in believing (or know). Differences in context do not make a difference, though differences in evidence or psychological state may. What is unclear is whether differences in what is at stake in a practical sense can raise the standards of evidence so that justification, and knowledge, can be

lost, or gained, with no change in evidence. But no one should be led to deny invariantism because of skepticism, for no good reasons have been offered for skepticism.

ACKNOWLEDGMENTS

I want to thank Elke Brendel, Mylan Engel, Mark Huston, Michael McKinsey, and John Turri for their comments on my paper. It was through discussion with Michael McKinsey that I came to see that fallacious skeptical arguments do not alter the evidence for *P*, even though, in a sense, they justify us in believing that *S* does not know *P*.

NOTES

- ¹ This is based on DeRose (2002, pp. 167–203 at 168–169).
- ² Jason Stanley gives these examples in Stanley (2004, pp. 2–3). He takes the first example from DeRose (1992, pp. 913–929 at 913).
- ³ DeRose prefers to put things metalinguistically, that is, he thinks the contextualist should be understood as saying that “*S* knows that *P*” is elliptical for “The sentence ‘*S* knows that *P*’ is true in context *c*,” not for “*S* knows that *P* in context *c*.” See, for instance, DeRose (1992, pp. 925–928). Applied to Hannah, he would say that “Hannah knows before the question is asked but not afterwards, even though her evidence and psychological state remain the same” sounds contradictory, but “The sentence ‘Hannah knows the bank will be open’ is true before she is asked the question,” but “The sentence ‘Hannah knows the bank will be open’ is false after she is asked the question about a possible change in hours (even though her evidence and psychological state remain the same)” does not sound contradictory. For me, these claims are on a par with respect to sounding contradictory: at first neither sounds contradictory and on reflection both do. However, nothing in my argument turns on putting things in the object language rather than metalinguistically. My argument will be that Hannah still knows after the question is raised by her husband and the sentence “Hannah knows that the bank will be open” remains true even after the question is raised by her husband.
- ⁴ In Russell (2001, pp. 34–48), I distinguish subjective from objective justification. The former involves epistemic blamelessness; the latter, support by the evidence. I argue that knowledge requires both sorts of justification.
- ⁵ Engel (2004). His discussion requires him to write of a *handless* BIV, but the argument is essentially the same.
- ⁶ Of course, if it were true that if *S* knows that *P*, then *S* knows that she knows that *P* (the KK principle), then insofar as she is justified in believing that she does not know that *P*, or even just believes that she does not know that *P*, she does not know that she knows that *P*. If the KK principle were true, it would follow that she does not even know that *P*. However, I do not believe that the KK principle is true. But perhaps a weaker principle is true, say, that if you know *P*, then you are

not justified in believing that you do not know that *P*. If that were true, then I would say that the evidential situation of Hannah and the student has changed once the implicit inference has been made. Then neither would have knowledge, but not for the reason the contextualist gives. They would lack knowledge because they would have an overriding (as opposed to undercutting) defeater.

⁷ Plantinga (1988, pp. 1–50 at 15–16 and note 18).

⁸ This suggestion was given on behalf of the contextualists by my colleague, Lawrence Powers, in the discussion following Jason Stanley's paper referred to in note 2. I have also been influenced by Powers' well-developed views about the central role of equivocation in philosophy.

REFERENCES

- Cohen, S.: 2001, 'Contextualism Defended: Comments on Richard Feldman's Skeptical Problems, Contextualist Solutions', *Philosophical Studies* **103**, 87–98.
- DeRose, K.: 1992, 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research* **52**, 913–929.
- DeRose, K.: 1999, 'Contextualism and Skepticism', *Philosophical Perspectives, Epistemology* **13**, 91–114.
- DeRose, K.: 2002, 'Assertion, Knowledge, and Content', *The Philosophical Review* **111**, 167–200.
- Engel, M. Jr.: 2004, 'What's Wrong With Contextualism, and a Noncontextualist Resolution of the Skeptical Paradox', *Erkenntnis* **61**, 203–231.
- Feldman, R.: 1999, 'Contextualism and Skepticism', *Philosophical Perspectives, Epistemology* **13**, 91–114.
- Feldman, R.: 2001, 'Skeptical Problems, Contextualist Solutions', *Philosophical Studies* **103**, 61–85.
- Plantinga, A.: 1988, 'Positive Epistemic Status and Proper Function', *Philosophical Perspectives, Epistemology* **2**, 1–50.
- Plantinga, A.: 1996, 'Respondeo', in J. Kvanvig (ed.), *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge*, Rowman & Littlefield, Lanham, 307–378.
- Russell, B.: 2001, 'Epistemic and Moral Duty', in M. Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, Oxford University Press, Oxford, 34–48.
- Scharifi, G.: 2004, 'Contextualism and the Skeptic', *Erkenntnis* **61**, 233–244.
- Sosa, E.: 1996, 'Postscript to "Proper Functionalism and Virtue Epistemology"', in J. Kvanvig (ed.), *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge*, Rowman & Littlefield, Lanham, pp. 271–280.
- Stanley, J.: 2004, 'Context, Interest-Relativity, and Knowledge', delivered March 3, 2004 at Wayne State University.

Department of Philosophy
Wayne State University
51 W. Warren
Detroit, MI 48201
USA
E-mail: bruce.russell@wayne.edu

WAYNE A. DAVIS

ARE KNOWLEDGE CLAIMS INDEXICAL?

ABSTRACT. David Lewis, Stewart Cohen, and Keith DeRose have proposed that sentences of the form “*S* knows *P*” are indexical, and therefore differ in truth value from one context to another.¹ On their *indexical contextualism*, the truth value of “*S* knows *P*” is determined by whether *S* meets the epistemic standards of the speaker’s context. I will not be concerned with *relational* forms of contextualism, according to which the truth value of “*S* knows *P*” is determined by the standards of the subject *S*’s context, regardless of the standards applying to the speaker making the knowledge claim. Relational contextualism is a form of normative relativism. Indexical contextualism is a semantic theory. When the subject is the speaker, as when “*S*” is the first person pronoun “*I*,” the two forms of contextualism coincide. But otherwise, they diverge. I critically examine the principal arguments for indexicalism, detail linguistic evidence against it, and suggest a pragmatic alternative.

1. INDEXICAL CONTEXTUALISM

The principal claim made for indexicalism is that it enables us to explain what Cohen calls *skeptical paradoxes*.^{1,2} Consider a proposition *O* we would ordinarily say we know, such as that we have a hand.

- A1 I know that I have a hand.
 $K(O)$

Since I can see my hand plainly before me, and feel it, what could be more evident? Now it is self-evident that if I have a hand, then I am not a brain in a vat. So it also seems evident that if I know that I have a hand, then I know that I am not a brain in a vat. The claim that I am a brain in a vat – one that receives all the sensory inputs I receive despite having no body – is a familiar skeptical hypothesis.

- B1 If I know that I have a hand, then I know that I am not a
 brain in a vat.³
 $K(O) \rightarrow K(\sim H)$

These two claims entail that I know that the skeptical hypothesis is false. But how could I, the usual argument goes, given that I would



have all the evidence I have now if I were a brain in a vat? This reasoning may lead us to deny that we know that the skeptical hypothesis is false.

- C1 I do not know that I am not a brain in a vat.
 $\sim K(\sim H)$

Since these three propositions are logically inconsistent, we have to give one up. But we have a hard time figuring out which. We may vacillate. What Cohen, DeRose, and Lewis want to explain is why we find this to be a difficult problem. Their answer is that we find “I know that I have a hand” acceptable because in ordinary contexts the relevant standard of evidence is low. We are entitled to ignore many alternatives. But mention of the skeptical hypothesis in the second proposition shifts us into a new context in which the standard of evidence is high. We accept the third proposition when we are in this elevated context. We find the skeptical paradox difficult to resolve, on their view, because of an unnoticed context shift.

It is obvious that “*S* knows *P*” has one form of indexicality, encoded grammatically in its tense. There is no other form of indexicality on the surface, but the indexicalist claims that the truth conditions of “*S* knows *P*” also vary with *the contextually indicated standard of evidence*. As Cohen points out, this semantic theory nicely explains why Smith, having just looked at his flight itinerary, would unhesitatingly say “Yes, I do” when asked if he knows whether the flight from Los Angeles to New York stops in Chicago. And yet Smith would say “No” later without embarrassment when Mary explains how critical it is that she gets to Chicago, and asks whether he really knows that the plane stops there. In the first context, Smith assumed the standard of evidence to be relatively low, while in the second he took the standard to be quite high. The fact that printed itineraries are occasionally wrong was relevant in the second context but not the first. The fact that our use of “know” is variable in this way is the solid kernel of truth in the contextualist theory.

Lewis’s theory of semantic dynamics, adopted by Cohen and DeRose, is that the mere suggestion of the skeptical hypothesis “changes the conversational score,” shifting us into a context with higher epistemic standards. The same thing happens when Mary points out how important it is that she gets to Chicago. These acts make certain possibilities of error “salient,” leading speakers to intend stricter standards.

2. DEROSE'S THEORY

DeRose (1995, pp. 22–23) observes that *explicit* skeptical claims do not produce a paradox. Let *O* be as before, but let *H* be: *My belief that I have a hand is false*.

- A2 I know that I have a hand.
- B2 If I know that I have a hand then I know that my belief that I have a hand is not false.
- C2 I do not know that my belief that I have a hand is not false.

We do not feel any paradox here. Once we have digested the double negative in C2, we unhesitatingly accept A2 and reject C2. Occasionally we may be in a skeptical mood, and question whether the evidence of our senses is sufficient for knowledge. In that mood, we accept C2 and reject A2. In neither case are we tempted to accept both A2 and C2. Since the suggestion of the skeptical hypothesis makes certain possibilities of error salient in the second example without producing the effects we are trying to explain in the first example, we cannot attribute the paradox presented by the first example to those factors. DeRose therefore refines Cohen's task: we have to explain why implicit skeptical hypotheses produce a paradox, but explicit skeptical claims do not.

Modifying the "tracking" theory of Dretske (1970) and Nozick (1981), DeRose (1995, p. 18, 27) describes *S*'s belief *P* as *sensitive* provided that *S* would not believe *P* if *P* were not true. Thus my belief that I have a hand is quite sensitive: if I did not have a hand, I would most certainly know it. I would see my arm end in a stump, I would not be able to wiggle my fingers, and there might well be considerable pain. But my belief that I am not a brain in a vat is insensitive. Even if I were a brain in a vat, I would still believe that I am not. For by hypothesis, the brain in the vat gets all the same sensory input my brain gets. Dretske and Nozick claim that knowledge entails sensitive belief. If it did, then we would have to conclude that I know that I have a hand (that belief is sensitive), but do not know that I am not a handless brain in a vat (that belief is insensitive).⁴ They concede that this result is counterintuitive, but accept it anyway, and thus reject the second, conditional proposition in the inconsistent triad.

DeRose finds conjunctions like (1) "abominable," and tries to avoid them by rejecting the claim that knowledge *always* requires sensitive belief.

- (1) I know that I have a hand but do not know that I am not a brain in a vat.

DeRose postulates a conversational “Rule of Sensitivity” whereby certain knowledge claims select a contextual standard for knowledge that does require sensitive belief. What DeRose claims, I believe, is (i) “*S* knows *P*” is true in context *C* provided “*S* believes – *P*” is true in every –*P* world that is epistemically relevant world in *C*; and (ii) the set of worlds epistemically relevant in *C* includes the nearest worlds in which *any* proposition mentioned or thought of in *C* as a possible object of knowledge is false.⁵ In ordinary contexts, skeptical hypotheses do not arise. But once someone asks whether I know I am not a brain in a vat, the set of relevant worlds has to expand to include some worlds in which I am brain in a vat. “I know $\sim H$ ” will be false in such a context because “I believe *H*” will not be true in any of the *H* worlds in this set. “I know *O*” will similarly be false because “I believe $\sim O$ ” will be false in those $\sim O$ worlds in the set that are *H* worlds. On this interpretation, DeRose avoids abominable conjunctions much the same way Cohen does.

When skeptical hypotheses have not been considered, the set of epistemically relevant worlds will be small. It will contain the closest $\sim O$ worlds, and “*S* believes $\sim O$ ” will be true in all such worlds. So *K*(*O*) comes out true. But since skeptical hypotheses have not been mentioned, the set of epistemically relevant worlds will not contain any *H* worlds. Since *S* will correctly believe $\sim H$ in all epistemically relevant worlds, *K*($\sim H$) comes out true too. So DeRose’s theory yields the desired result that both *K*(*O*) and *K*($\sim H$) are true in everyday contexts.

DeRose’s Rule of Sensitivity, however, produces a very curious result. It implies that while “*S* knows $\sim H$ ” is true in everyday contexts, we cannot *say* or even *think* that *S* knows $\sim H$ in everyday contexts (DeRose, 1995, p. 39).⁶ For saying or thinking that *S* knows $\sim H$ would raise epistemic standards by including some *H* worlds in the set of relevant worlds. Since “*S* believes $\sim H$ ” would be true even in such worlds, “*S* knows $\sim H$ ” would now be false. As far as I am concerned, the implication that it is impossible for us to truly assert an instance of “*S* knows *P*” – *even when it is true* – is at least as abominable as the Nozickian conjunction. “*S* knows $\sim H$ ” does not seem at all like other sentences that can be true as long as they are not used assertively, such as “I am not asserting anything.”

3. FAILURE TO EXPLAIN THE PARADOX

While even opponents credit the indexical theory with providing a plausible explanation of why we find certain inconsistent triads paradoxical, the appearance of success is illusory. A defining characteristic of the theory is the claim that the truth value of “*S* knows *P*” is determined by the standards prevailing in *the context of attribution*, not by those in *the context of the subject*. When we are in a high standard context, we have to deny that anyone knows *O*, even subjects in low standard contexts; in particular, we have to deny that we ever know *O*.⁷ DeRose does raise this question.

[W]hy do we find these claims to know plausible even when we’re in a context in which the skeptic has raised the standards to such a level that these claims are false? (DeRose, 1995, p. 40)

But his answer misses the point.

[W]e at the same time realize that as soon as we find ourselves in more ordinary conversational contexts, it will ... be true for us to claim to know these very *O*s that the skeptic now denies we know ... (DeRose, 1995, p. 41)⁸

The paradox to be explained is constituted by the fact that it seems to us that we know we have a hand even in contexts – like our context now – in which we have considered the skeptical possibility that we are brains in vats, which seems to imply that we do not know any such thing. We want to deny the skeptical hypothesis, but do not know how we can. We are unsure whether to give up $K(O)$, $\sim K(\sim H)$, or $K(O) \rightarrow K(\sim H)$. Many believe all three propositions, while realizing that they are incompatible. If knowledge claims were indexical, we should clearly and willingly deny $K(O)$ after standards have been raised by considering $K(\sim H)$. We should regard the fact that we will accept $K(O)$ when we are in an ordinary context as completely irrelevant to whether $K(O)$ is true in our current context, just as we regard the fact that we accept “I am a foreigner” when I use it in Germany as irrelevant to whether it is true when I use it in America.

Suppose we somehow resist any upward pressure on standards when we consider $K(\sim H)$, and remain in a low standard context, as DeRose and Cohen both allow. Then we should unhesitatingly accept $K(O)$ and reject $\sim K(\sim H)$. The contextualist theory predicts no paradox in this case too. There should not even be a paradox if the standards shift up and down without us knowing it. We might want to explain why we sometimes confidently accept $K(O)$ and $K(\sim H)$ and other times just as confidently reject them. But in neither condition

will we be making any hard choices or believing incompatible propositions.

DeRose noted that Cohen's theory failed to discriminate between the paradoxical and the non-paradoxical triads. DeRose sought to explain what is distinctive of paradoxical triads by adopting a more externalist view of standards. But that does not help. Once we consider whether I know that *my belief that I have a hand is not false*, the Rule of Sensitivity says that the set of relevant alternatives must include worlds in which my belief that I have a hand *is* false. Since I still believe I have a hand in such a world, the Rule of Sensitivity would force us to give up $K(O)$ for the same reason it forces us to give up $K(O)$ in the paradoxical case.⁹

If the hypothesized indexicality of "*S* knows *P*" did explain why we find it difficult to resolve the paradoxical triad, then we should expect to find similar paradoxes with other indexical expressions. Consider:

- A3 Edward's statement that Henri is a foreigner was true.
- B3 If Edward's statement that Henri is a foreigner was true, then Monique's statement that Henri is a foreigner was true.
- C3 Monique's statement that Henri is a foreigner was not true.

We do not find any paradox here once we know that Edward is English and Monique French. We instantly reject the conditional proposition B3. Judging from this case and the others to be presented, we *should* embrace the abominable conjunction if knowledge claims were indexical. Neither Cohen nor DeRose want to do that.

Another example is less familiar but more relevant. A result is said to be *significant* in statistics if the likelihood of its being due to mere chance is close enough to zero to be confident that it was not a chance result. As is well known, there is no non-arbitrary way to specify how close is close enough. The conventional choices are $p < 0.05$ or $p < 0.01$, but others are possible. When an unqualified statement like "The results are statistically significant" is made, there is an implicit reference to a contextually indicated significance level. We have no trouble seeing how to resolve the apparent inconsistency in the following three statements:

- A4 Edward's statement that the result is significant was true.
- B4 If Edward's statement that the result is significant was true, then Monique's statement that it is insignificant was false.
- C4 Monique's statement that the result is insignificant was true.

Suppose there is a 0.03 likelihood the experimental results were due to chance. If we know that in Edward's context the significance level was 0.05 while in Monique's it was 0.01, then we instantly accept the first and third statements and again reject the conditional. But if we know that in both contexts the significance level is 0.01, then we keep the conditional and reject the first statement. Neither case provides a good model for skeptical paradoxes. Of course, it is possible that "*S* knows *P*" has a *sui generis* form of indexicality. But the fact that uncontroversially indexical expressions do not generate similar paradoxes casts some doubt on the claim that skeptical paradoxes result from the indexicality of knowledge claims.

To explain skeptical paradoxes, the indexicalist must explain why we continue to want to assert $K(O)$ (or deny it reluctantly) even when we have shifted into a high-standard context in which it should be obviously false. One possible explanation is that "The apparent closure failures are illusions that result from *inattention to contextual shifts*" (Cohen, 1988, p. 111, my emphasis).¹⁰ But if we still think low everyday standards prevail, we should unhesitatingly affirm "I know that I am not a brain in a vat" (C1). If we do not notice that we are in a high standard context, then the indexicalist cannot say that we are inclined to reject $K(\sim H)$ because consideration of the skeptical hypothesis shifts us into a high standard context. The inattention thesis also undermines the predictions that the indexical theory seems to have gotten right. The example of Smith at the airport is taken as independent evidence that "know" is indexical. But this interpretation of the example is plausible only if Smith notices the change in context that results when Mary stresses how important it is that she get to Chicago. If competent speakers do not attend to contextual shifts in their use of "know," what reason is there to say that the term is indexical? If it is suggested that competent speakers generally notice relevant context shifts in other cases, then the indexicalist needs to explain what is special about cases in which implicit skeptical hypotheses are considered that induces inattention. Without such a specification, the indexicalist has not explained the skeptical paradoxes.

Cohen later suggested that speakers are "*unaware of, and so misled by, the kind of context-sensitivity ... involved in ascriptions of knowledge*" (1999, p. 79, my emphasis). Our use of "*S* knows *P*" varies with contextually indicated standards, but we do not realize this; we are "blind" to its indexicality. Schiffer charged that this "error theory" refutes Cohen's semantics: "Speakers would know what they were saying if knowledge sentences were indexical in the

way the Contextualist requires” (1996, p. 328). The fact that we are not blind to the indexicality of other words supports Schiffer’s objection.

Cohen (1999, p. 61) responds by observing that similar blindness is possible with “flat.”

[T]here is nothing implausible about combining a contextualist semantics with an error theory. Consider ascriptions of flatness. You can lead competent speakers to question their everyday ascriptions of flatness by making salient “bumps” that ordinarily we do not pay attention to. Taking this strategy to the extreme – e.g., by calling attention to microscopic irregularities – one can lead competent speakers to worry about whether anything is really flat. (Cohen, 1999, p. 78)

Cohen cites Unger (1975) as proof of this possibility. However, it is one thing to claim that competent speakers can be led to form erroneous beliefs about the meanings of common terms, and another to claim that such an error is common. In fact, competent speakers routinely apply the term “flat” to roads, fields, and table-tops they know have some bumps in them – things they would readily describe as *flat but of course not perfectly flat* without any hint of contradiction. Competent speakers can choose a standard of flatness so high that even microscopic irregularities count. Alternatively, they can use “flat” in its geometric sense. They could then say truly “Nothing is flat.” A man in either mood might deny that “flat” could ever be properly applied to Iowa. He would be making an egocentric mistake. While some competent speakers might refuse to recognize any other use of “flat,” most will readily grant that we can correctly say that Iowa is flat compared to Colorado and hilly compared to Florida, and that what we mean when we say simply “Iowa is flat” depends on what we are comparing it to as does what we mean when we say “Iowa is big” (Texas, or Rhode Island?).

Furthermore, we have no difficulty resolving the apparent inconsistency in the following triad:

- A5 Our claim (after spending a year in Colorado) that Iowa is flat was true.
- C5 Our claim (after spending a year in Florida) that Iowa is hilly was true.
- B5 If our claim that Iowa is flat was true, then our claim that Iowa is hilly was false.

We recognize that when making our first claim (A5), we were comparing Iowa to Colorado, whereas we were comparing it to Florida when making our second claim (C5). The different standards of comparison allow both to be true and make B5 false. If we insist that

all claims should be evaluated using a strict standard of flatness, then we will accept C5 and B5 and reject A5. If we insist on loose standards, then we will accept A5 and B5 and reject C5. So even people who are blind to the indexicality of “flat” will find no paradox like the one we are trying to explain in the case of knowledge. We get the same result with “foreign” and “statistically significant.” Cohen’s error theory thus undermines not only his semantics, but the explanation he proposes for the skeptical paradoxes. If speakers do not realize that “know” is context sensitive, then there is no reason to expect their use of “know” to vary with contextually indicated standards.

4. FAILURE TO EXPLAIN THE LOTTERY DIFFERENCE

Cohen (1988, pp. 106–109) believes the contextualist theory explains another puzzling feature of our commonsense knowledge attributions. It seems evident that no matter how many people have entered a lottery, that fact alone does not enable us to know that we will lose. We commonly grant, though, that we can know that we lost based on a newspaper report listing someone else as the winner. Yet we know that the likelihood of an error in the newspaper may be much greater than the likelihood that we will win the lottery. What is the difference? Cohen’s explanation invokes Lewis’s theory of salience as a determinant of indexical reference.

[T]he contexts in which we evaluate the two cases are different, i.e., the standards of relevance that operate are different. The explanation for why the standards differ lies in the fact that in the statistical case, unlike the other cases, the nature of the reasons make the chance of error salient. (Cohen, 1988, p. 107)

But suppose we are given that Ed believes Tom lost because the odds of winning are one in a million, while Bill believes Tom lost after reading about the winner in the paper. *We* are inclined to say that Bill knows Tom lost, but Ed does not. But *we* are in a context in which, if Cohen’s hypothesis is correct, the standard of evidence has been raised by the salience of the one in a million chance of error. So once that high standard of evidence has been contextually selected, Cohen’s hypothesis implies that we should reject the claim that Bill knows Tom lost as well as the claim that Ed does. Yet we do not. Cohen correctly observes that contexts may shift back and forth. But my example involves judging *both* “Bill knows Tom lost” *and* “Ed does not know Tom lost” in the *same* context. No matter how small

the chances of being wrong about losing the lottery are, we are not willing to say that we know that we will lose even with the “everyday” standards that allow us to say that we know based on testimony and perception.¹¹

5. EVIDENCE AGAINST INDEXICALITY

I have already provided one reason to doubt that “*S* knows *P*” is indexical in the way the contextualist theorizes: we should not find inconsistent skeptical triads to be paradoxical if “*S* knows *P*” were indexical. This was shown both by examining what truth values the theory would assign and why, and by examining how known indexicals behave in similar contexts. The supplementary hypothesis that competent speakers are blind to the indexicality of knowledge, or fail to respond properly to contextual shifts, does not help either, and the latter undermines the independent evidence cited for the theory.

A related piece of evidence is this. When Edward says “The result was significant” and Monique says “The result was insignificant,” they appear to be contradicting each other. But there is no contradiction if Edward and Monique have chosen different significance levels. In that case we can say “They are both right.” Similarly, when the Coloradan says “Iowa is flat” and the Floridian says “Iowa is not flat,” they are both right because they are comparing Iowa to different states, just as the Texan and the Rhode Islander are both right when one denies and the other affirms “Pennsylvania is big.” But when Moore, focusing on the evidence of his senses, proclaims “I know that I have a hand,” and the skeptic, focusing on various remote possibilities proclaims “No one knows that he has a hand,” the last thing we are inclined to say is “They are both right.” Similarly, Smith’s second answer “No, I don’t” seems to clearly *contradict* his first answer “Yes, I do” in the airport case. Smith’s second answer seems to be a *concession* – a *retraction* or *qualification* of his earlier claim upon further consideration of the strength of his evidence. And when we come to deny that we have knowledge under skeptical pressure, we regard our earlier claims to have knowledge as *mistakes*.¹² The claim that competent speakers find such utterances contradictory because they are blind to the indexicality of “*S* knows *P*” seems completely *ad hoc*.

In the significance case, the standard can go down as easily as up. If we learn that the level of significance mutually accepted in

the conversation is 0.05 rather than 0.01, then we instantly accept the first two statements in the triad and reject the third. DeRose noted, however, that the contextual shifts in the epistemic cases have a tendency to go up that is unexpected given the indexical theory.

Why can't the commonsensical epistemologist simply declare again that he knows, and rely on a Rule of Accommodation to lower the standards back down so as to make *his* claim true? To this Lewis responds that, for some admittedly unknown reason, the standards are more easily raised than lowered (Lewis, 1979, p. 247) (DeRose, 1995, p. 8, fn. 12)

Cohen has suggested that geometric terms also displays downward stickiness. If I say "France is hexagonal," someone might raise the standard by saying, "No, look at the way Normandy sticks out." Cohen observes that we cannot lower the standard simply by repeating "France is hexagonal." This casts doubt on Lewis's Rule of Accommodation, perhaps. But we *can* easily lower the standard, by saying "I'm just concerned with its rough overall shape."

Compare the following arguments.

- (2) *S* spoke truly when saying "I am not a man."
Therefore, I am not a man.
- (3) *S* spoke truly when saying "Wayne Davis is a foreigner."
Therefore, Wayne Davis is a foreigner.
- (4) *S* spoke truly when saying "Some people know that they have a hand."
Therefore, some people know that they have a hand.

Because they involve indexicals, the first two arguments are patently invalid. But the third strikes us as valid. Anyone who seriously offered the first argument would prove that he or she did not fully understand the first-person pronoun. Anyone who thought the second was valid must either not understand "foreigner" or not realize that *S* might be a foreigner. In contrast, many competent speakers of English have thought the third argument was valid. If it is not, the defect is not obvious.¹³

Similar evidence emerges by reflecting upon the following caveat.

[S]o strictly speaking, instead of saying that *S* knows *P* in one context but fails to know in another, one should really say that the sentence '*S* knows *P*' is true in one context and false in another. Because these metalinguistic locutions are stylistically cumbersome, I will continue to speak instead in the object language. (Cohen, 1999, p. 65)

Consider the following:

- (5) a. "The Hope diamond looks blue" is true in one context,
but not in others.
- b. The Hope diamond looks blue in one context, but not in
others.

Strictly speaking, these are different statements. It is even possible for the first to be false when the second is true, because the sentence mentioned in (5a) might have a different meaning. But given the meaning that the sentence mentioned in (5a) actually has, these two statements are non-logically equivalent: one is true if the other is. Given this equivalence, we can let stylistic considerations dictate our choice of words. When there is indexicality beyond what is contained in present tense verbs, there is no such equivalence. Compare:

- (6) a. "He is my son" is true in one context, but not in others.
- b. He is my son in one context, but not in others.

The two statements in (6) are not even close to being equivalent. The first is true, the second false. On Cohen's view, (7) should be more like (6) than (5).

- (7) a. "*S* knows *P*" is true in one context, but not others.
- b. *S* knows *P* in one context, but not others.

But in fact, (7) seems more like the non-indexical (5).¹⁴

While Cohen cites "flat" as a model for the indexicality of "know," we have observed some crucial differences. Another is that we apply the adverbs *completely* or *perfectly* to *flat* when we want to adopt the strictest possible standard. A surface is perfectly flat when it has *absolutely no* curvature or unevenness. Florida is far from perfectly flat. Indeed, we recognize that nothing in the material world is perfectly flat, even the most precisely milled surface. "Completely" and "perfectly" cannot modify "know." There is no contrast between perfect and imperfect knowledge.¹⁵ A related difference between "know" and "flat" is that only the latter allows quantitative comparison (cf. Dretske 1981, p. 363). We know that Florida is *flatter* than Colorado, but it makes no sense to say that one person's knowledge that he has a hand is greater or lesser than another person's. Similarly, only "flat" can be modified by *compared to X*. Unlike "perfectly flat" or "flat compared to Colorado," "flat" *simpliciter* is a relative term: something is flat provided that it is at least as flat as the contextually indicated standard of comparison. It must have no "significant" curvature or unevenness.

A final and ironic piece of evidence against indexicalism is provided by the fact that it does not account well for the meaning of skepticism. Consider a standard formulation:

- (8) No one knows anything.

This sentence expresses a controversial philosophical thesis, one difficult to prove false without begging the question. But suppose indexicalism were correct. Then (8) would express either an obvious truth (in a skeptical context) or an obvious falsehood (in an everyday context). Even if the auxiliary blindness hypothesis is added, saying that no one notices that (8) expresses different propositions in different contexts, indexicalism still predicts that (8) never expresses a controversial thesis. Contrast “flatness skepticism.”

- (9) Nothing is flat.

If “flat” is used according to geometric standards, then (9) will seem obviously true. It will seem obviously false if it is used according to Coloradan, Floridian, or any other terrestrial standard. There is nothing controversial or puzzling on either interpretation. If anyone were blind to the indexicality of “flat,” they might be puzzled as to why they sometimes affirm (9) and sometimes deny it. But at no time will they find themselves puzzled as to whether what (9) claims is true.

Another way of formulating skepticism is provided by (10).

- (10) People ordinarily believe that they know things, but they do not.

If indexicalism were correct, this sentence should not express a controversial thesis either, but for different reasons. If “know” were being used according to skeptical standards, then the first conjunct of (10) is clearly false (people ordinarily do not believe they know things according to the elevated standards of skeptical contexts). If “know” were being used according to ordinary standards, then the second conjunct of (10) is clearly false. If we equivocate, then (10) will seem to express either an uncontroversial truth (first conjunct low, second high) or an uncontroversial falsehood (first high, second low). In no case would (10) seem to express anything controversial if indexicalism were true. This remains true even if we add the auxiliary hypothesis that people are blind to the indexicality of “know.” The indexicalist attempt to explain why skeptical hypotheses generate paradoxes must be inadequate if the theory predicts that skepticism is not even a controversial thesis.

6. NORMATIVITY

Perhaps the most important feature of “*S* knows *P*” that the indexical theory ignores can be seen by reflecting further on the explanatory failure of the indexical theory. If the indexical theory were correct, we should not have difficulty figuring out which proposition in the incompatible epistemic triads to reject, any more than we have difficulty in the triad in which significance levels are relevant. Why then *do* we have trouble? I think the reason is simply that consideration of the skeptical hypothesis makes us reconsider whether the standard of evidence we ordinarily accept is the *proper* standard.¹⁶ I believe firmly, for example, that the evidence of my senses is not only the best evidence I can have that I have a hand, but all the evidence I need to know this. Yet the fact that I would have all the same evidence if I were a brain in a vat makes me wonder whether I do not need some more evidence to rule that possibility out. I convince myself that more evidence is unnecessary by reflecting on the fact that the brain in the vat hypothesis is completely *ad hoc*; that we have never observed anything like a disembodied brain getting all the same sensory stimulation as a real human, and cannot realistically see how anything like it could be arranged; that it is more complex than the simple hypothesis that I have a hand; and so on. Still, the hypothesis is a logical possibility. I recognize that others are unable to convince themselves that more evidence is unnecessary, and as a result have a skeptical moment. I remember having had one myself when first studying epistemology, and still wonder occasionally about the strength of my own reasons.

I believe we vacillate over which proposition in the inconsistent epistemic triad to give up because we have no obvious way of *settling* the question as to what the proper standards for knowledge are. It is not clear how we can proceed without begging the question. The skeptical paradoxes raise foundational issues. The most serious flaw in the indexical theory of knowledge is that it makes any such normative deliberation as otiose as a debate over whether the 0.01 significance level is more proper than the 0.05 or 0.10 level. Within limits, the choice of a precise significance level is up to the investigator. The choice of epistemic standards is not. Similarly, it makes no sense to deliberate about whether the Floridian’s standard or the Coloradan’s is more appropriate for judgments of flatness. Since no standard of comparison is intrinsically correct, we are free to choose. No deep issues are raised. When the skeptic questions whether we really know the things we think we know, suggesting that the evi-

dence of our senses is not good enough, it is not adequate to say that we have a different standard of evidence in mind. The question is whether the standard we have in mind is correct.

7. AN ALTERNATIVE THEORY

It is undeniable that our use of “*S* knows *P*” seems to presuppose different epistemic standards in different contexts, suggesting that it is indexical. We have seen, however, that the indexical theory does not account well for the observed variability. To account for some of the facts, moreover, the theory needs to postulate that competent speakers are ignorant of the indexicality of knowledge claims, or inattentive to the relevant context shifts. In addition to being implausible if “know” has an indexical meaning, the auxiliary error thesis undermines the explanation of other cases. Are there alternative explanations? I believe that the observed variability has at least two sources, both pragmatic.

In at least some cases, the difference in what we assert or accept can be attributed to a difference in what we consider to be the proper standard for a particular case. The pressure of the skeptical considerations makes us *change our minds*. We may assert “I know that I have a hand” before thinking about skeptical hypotheses, and deny it at least temporarily afterward. The fact that philosophers like Hume deny they know things when doing philosophy that they readily claim to know in everyday life may simply be due to the fact that they change their minds about the requirements of knowledge when they do philosophy. In the same way, some people vacillate over the existence of God because sometimes “faith” prevails, and other times “reason.” The fact that they assert “God exists” in Church but deny it in the seminar room does not imply that “God exists” is indexical.

But Cohen’s example of Smith and the stop in Chicago does seem to show that our use of “know” is sensitive to contextual factors. Recall that Smith said “Yes” when asked if he knows whether the flight stops in Chicago, while later saying “No” after learning how important it is that Mary get there. This difference in what Smith says does not seem attributable to the fact that Smith changed his mind about whether a flight itinerary is a sufficient basis for claiming knowledge. Nor does it seem attributable to his suddenly remembering that the itinerary might be wrong. In airports, these things are usually on our minds. When he said “No,” Smith is unlikely to have

felt embarrassed about what he said earlier, or to have regretted saying it. He probably did not have the wrenching feeling that he knows less than he thought he did, as we do when we are influenced by skeptical hypotheses.

So what is going on when we change what we say we know without changing our mind about what we know? One plausible hypothesis is that Smith was interpreting “I know” in the first context *loosely*, to mean or imply that *he is sufficiently close to knowing for the purposes at hand*. When the context required greater precision, he stopped speaking loosely. The same account is plausible for Vogel’s (1999, p. 161) car example. If we are walking out to the parking lot, and you ask “Do you know where your car is?” I will ordinarily answer “Yes” as long as I remember where I left it. But if you go on to ask whether I am sure no one stole or towed my car, and then ask again whether I know, I will answer “No.” In the first instance, I would ordinarily interpret you as wanting to know whether I remember where I parked my car. I would assume you were not interested in the well known but unlikely possibilities of theft or towing. So “Yes, I know” was an appropriate answer because I was close enough to knowing for the presumed purpose of the conversation. But once you have mentioned the possibility of theft or towing, I would realize that I am not close enough to knowing given your interests, and would give a negative answer.

We exhibit similar variability in our use of other terms, such as “all”. If a lawn service uses a chemical that kills 99% of our grass, we are likely to report that to the company by saying “You killed all the grass”. In this case, what we said is literally false.¹⁷ But what we implicated is true, namely, that the percentage of the lawn killed by the chemical was sufficiently close to one hundred for the purposes of the conversation.

Meaning or implying something that we do not say is called *implicature*.¹⁸ *Conversational implicatures* are those in which what the speaker means or implies is not part of what the words used mean, but depends instead on the conversational context. So what we mean or imply when we use “know” or “all” loosely is a conversational implicature.¹⁹

Loose speech is similar to hyperbole, in which we say something for effect that we know to be too strong. In loose speech, we need not believe that what we said is literally false. Indeed, we need not have taken the time to check, as is often the case when “all” is used after a quick glance at the set of objects in question. We sometimes do use “know” and “all” hyperbolically, as when we say “I just

knew that he was going to screw up” or “He *always* screws up.” But when Smith said “Yes, I do,” he did not intend his audience to recognize that he did not literally know that the plane would stop in Chicago. Indeed, he may not have reflected long enough on the question to take a definite stand. Similarly, when we take one look at our decimated yard, and complain that the lawn company killed all the grass, we may not have investigated closely enough to determine whether literally all is dead, because we do not need to know more than we do. “Perfectly flat” is often used hyperbolically, mimicking the literal use of the relative term “flat.” When a carpenter proudly announces that he sanded the wood perfectly flat, he knows that what he says is not literally true. But what he is implicating is that he achieved a remarkably high degree of flatness, with no deviations large enough for even a trained eye or hand to detect.

Grice observed that conversational implicatures are generally *cancelable*. Thus a standard conversational implicature of “Some *S* are *P*” is “Not all *S* are *P*.” This implicature can be cancelled by adding “indeed, all are.” The implicature may also be cancelled by saying “Some *S* are *P*” in a context in which it is obvious that every single *S* is *P*, in which case it is an understatement, the opposite of hyperbole. However, implicatures are *not* cancelable when they are entailed by what is said. Thus if someone asks, “Have you been to Belgium or Luxemburg?” we can answer affirmatively by saying “I’ve been to Belgium.” The affirmative answer is implicated, since we implied but did not say that we have been to Belgium or Luxemburg. And the implicature is conversational, since it would not have been possible if there had been no mention of Luxemburg. But this implicature cannot be canceled without violating the laws of logic. The implicatures involved in loose use are non-cancelable for the same reason. “I know” entails “I am close enough to knowing for the purposes at hand,” just as “You killed all of the grass” entails “You killed all or close enough to all of the grass,” and “It is perfectly flat” entails “It has no significant bumps”.²⁰

In the case of “I know that I have a hand” – *K(O)*, the issue is whether it is literally true, not just loosely but strictly speaking. Unlike the airport case, we do not readily retract such a claim when someone asks “Do you know for sure?” Indeed, outside of philosophy we regard the question as silly or impertinent. When doing philosophy, the suggestion that we might be a brain in a vat sometimes shakes our confidence, and leads to a retraction. In my view, we should reject skeptical hypotheses and affirm the literal

truth of $K(O)$. The skeptic's standard are too strict even for the strict use of " S knows P ." Others disagree. This is a difference of opinion about epistemic norms, not a case of indexical variation. Cohen has claimed that the indexical theory's ability to provide a definitive resolution to this issue is a plus. But as we noted, the indexical theory predicts that $K(O)$ when used strictly should seem clearly false once skeptical hypotheses have been considered. I do not believe this prediction is borne out. Indeed, we set out to explain why it is hard to make up our minds in this case. I myself do not look with favor on a theory conceding that "I know I have a hand" is false in any context.

The hypothesis that " S knows P " has an indexical meaning is a *semantic* theory. The alternative I have sketched accounts for the contextual variability as a *pragmatic* phenomenon. Since this pragmatic account does not claim that " S knows P " is indexical, it avoids several of the difficulties with the contextualist theory. It allows that " S knows P " and " S does not know P " are contradictory; it allows that the proper epistemic standards are not determined by choice or intention; and it does not predict that the paradoxical epistemic triad should be as easy to resolve as the parallel indexical triads. It does not require attributing widespread linguistic error to competent speakers. The pragmatic theory is also able to explain why standards for the use of "know" are more readily tightened than relaxed. When we have spoken loosely, or offered an approximation, greater precision is often demanded and provided. But once we have spoken strictly, or made a precise claim, it is generally either pointless or confusing to begin speaking loosely. After I have reported that 95.74% of my grass was killed, our audience is unlikely to interpret "All my grass was killed" as loose speech. It will thus seem like a correction, or a report that things have gotten worse. Finally, my pragmatic alternative allows a better explanation of the skeptical paradoxes.

Neither of the two pragmatic sources of variability I cite accounts for the lottery difference. But the difference is not merely that we affirm " S knows on the basis of testimony" in some contexts and " S does not know on the basis of odds" in other contexts, but rather that we judge both in the same context. So an account that does not cite contextual differences is needed. One difference that may be relevant is that in the lottery case, we are certain that someone using exactly the same reason will end up with a false belief. We have no such certainty in the testimony case. This suggests that there is a contextually invariant requirement at work.

8. THE SYSTEMATIC FALSEHOOD ARGUMENT

DeRose says that the contextualist theory “is designed largely with the goal in mind of crediting most of our attributions of knowledge with truth” (1995, p. 46). I presume “most” was used very loosely here. I am sure DeRose would not want to defend the claim that over 50% of all actual uses of instances of “*S* knows *P*” were true. A more cautious formulation of his premise follows:

We in general take it as a strike against a theory of a common term of a natural language that it involves the speakers of that language in systematic and widespread falsehood in their use of that term. (DeRose, 1995, p. 46)²¹

The claim that contextualism avoids involving speakers in “systematic” falsity is questionable given the contextualist’s view that uses of “*S* knows *P*” are always false in contexts in which the skeptic’s epistemic standards are operative. If contextualism is correct, there should also be contexts in which the standards are high enough to make most but not all knowledge claims false. The claim that indexicalism avoids “widespread” falsehood is more plausible, given that many instances of “*S* knows *P*” come out true in everyday contexts. But by anyone’s standards, everyday knowledge claims are frequently erroneous. Think of all the instances in which *P* is about the uncertain future, the distant past, things misperceived, religious matters, or outdated scientific results; along with all the instances in which *S* is an ignoramus.

Cohen (1999, p. 83) is right that indexicalism enables us to avoid skepticism in everyday contexts. But it is not clear why this should be cited as an advantage on *linguistic* grounds. “Witch” is still a common term of English, and used to be used regularly in daily life. Nevertheless, its meaning is such that all instances of “*S* is a witch” are false. The same goes for “Santa Claus.” And consider the term “god.” Given that there are so many different religions making incompatible claims about gods, there is widespread and systematic falsehood in the use of “god” and its translations in other languages. Only those speakers of the true religion, if there is one, will speak truly. A wide range of knowledge attributions would have to be true if “knowledge” were an observation term like “water,” “cat,” or “thought,” expressing a concept acquired as a result of perceptual or introspective contact with its instances. But it is evident that “knowledge” is not observational.

Even if DeRose has identified a general presumption guiding semantic analyses, we would have to decide why it should take

precedence in the case of “knowledge” over other principles that pull in different directions. I noted above that we take it as a strike against a theory of a term learned early in life and used regularly that the theory postulates that competent speakers are ignorant of important features of the term’s meaning. Yet that is precisely what the contextualist does to explain the skeptical paradoxes. Why should we give more weight to avoiding attribution of falsehood than to avoiding attribution of linguistic mistakes or blindness?

Finally, let us grant for the sake of argument that speakers generally use “*S* knows *P*” to communicate truths. A separate argument is needed as to whether they are communicated as part of what is said rather than as something implicated. Suppose that speakers always use “know” to say something that is false while implicating something that is true. Are speakers engaged in “systematic and widespread falsehood in their use of the term”? DeRose’s principle is not precise enough to yield any definite conclusion in this case, and it is not easy to see how the principle could be strengthened so as to favor truth in what is said over truth in what is implicated.

DeRose argues for his methodological principle as follows.

[Suppose] that a crazed philosopher claimed that there are no physicians, because, in addition to holding a medical degree, a necessary condition for being a physician is that one be able to cure any conceivable illness. On what grounds should we reject this bizarre conjecture in favor of a more traditional and less demanding account of what it is to be a physician? ... it’s eminently reasonable to suppose that such facts as these, regarding our use, in thought and in speech, of the term ‘physician’ are involved: that we take to be physicians many licensed practitioners of medicine who don’t satisfy the demanding requirement alleged; that we seriously describe these people as being physicians; ... etc. (DeRose, 1995, p. 47)

The fact that competent speakers routinely apply “physician” to people they *take* to be far from satisfying the demanding requirements is indeed evidence that such requirements are not part of the meaning of the term. If instead they used the term “physician” precisely to convey the belief that those requirements are satisfied, then the objective fact that those people do not satisfy them would not undermine the conclusion that the requirements are part of the meaning of the term. That is, what “physician” means is determined by what *beliefs* or *thoughts* speakers use it to express, whether true or false.²² Similarly, given that speakers of English conventionally use “*S* knows *P*” to express thoughts and beliefs about what people know rather than about what people eat, we can draw conclusions about what “know” means in English. We do not need to know whether the

beliefs expressed are true. If someone proposes that part of what “*S* knows *P*” means is “*S* can rule out all conceivable alternatives to *P*,” we can gather evidence against this semantic hypothesis by observing that people who use “*S* knows *P*” commonly do not have such a belief, are not trying to indicate that they do, and may even agree that *S* cannot rule out all conceivable alternatives when the question is posed. If we found on the contrary that speakers intended to communicate the idea that *S* can rule out all conceivable alternatives, then we would have to take that as part of the meaning “*S* knows *P*” even if that meant that its instances are generally false. Whether a particular belief is commonly expressed by users of “*S* knows *P*” is relevant to what “know” conventionally means; the truth of the belief is irrelevant.

In response to Cohen, it should also be noted that the particular pragmatic theory I have offered does allow that “*S* knows *P*” is literally true in many everyday contexts. When terms are used loosely, the speaker is not committed to the falsity of what was literally said. We can concede that when Smith said he knows the plane makes a stop in Chicago, what he said was literally false, while insisting that when Moore says he knows he has a hand, what he said is literally true. The proper way to defend a theory of the meaning of “know” against the charge of widespread falsehood, I believe, is to provide arguments against skepticism.

ACKNOWLEDGMENTS

I would like to thank Jack Bender, Gerhard Ernst, Christoph Jäger, Gerhard Minnameier, Mark Siebel, Stewart Cohen, and an anonymous referee for many helpful comments on earlier drafts of this paper. I would also like to thank Elke Brendel and Christoph Jäger for organizing the outstanding conference “Contextualist Approaches to Epistemology” (University of Mainz, September 4–6, 2003), at which I presented the previous draft of this paper and learned a great deal.

NOTES

¹ See Lewis (1979, 1996), Cohen (1986, 1987, 1988, 1999, 2000a, and b), DeRose (1995). See also Unger (1975, 1984, pp. 46–54), Dretske (1981, p. 367, 376), Heller (1999), Klein (2000), Fogelin (2000a, b), Valdés-Villanueva (2000) and Rosenberg (2000).

- ² Austin (1961, p. 56) appears to suggest the relational form of contextualism, as does Annis (1978, p. 207), Dretske (1981, p. 377), Sosa (2000, p. 142), and perhaps Williams (2000, p. 82). See Feldman (1999, p. 94) for a clear statement of the distinction.
- ³ This conditional is often derived from a general principle of epistemic closure. Any generalization of B1 raises questions that B1 itself does not.
- ⁴ Cf. Cohen (1999, p. 63, 70) and Feldman (1999, pp. 98–99).
- ⁵ Cf. Heller (1999).
- ⁶ Cf. Cohen (1999, p. 71) and Lewis (1996, p. 420, 434, 441–442). DeRose's (1995, p. 39) most official formulation of the rule says that epistemic standards "tend to be raised, if need be" so as to require sensitive belief. But then his theory does not completely avoid abominable conjunctions, and needs to explain when shifts do occur. Cf. Cohen (1999, p. 81). Contrast Feldman (1999, pp. 100–101).
- ⁷ Cf. Schiffer (1996 pp. 324–325), Fogelin (2000b, p. 92) embraces this conclusion, but wrongly concludes that he is thereby rejecting the indexical theory.
- ⁸ Cf. Lewis (1996, pp. 441–444) and Cohen (1999, p. 67; 2000a, p. 102). Elke Brendel and Mark Lance pointed out another problem with this claim: our statement in high standards contexts that "*S* knows $\sim H$ " is true in ordinary contexts implies that " $\sim H$ " is true given the facticity of knowledge. But then indexicalism commits us to "*p* but I do not know that *p*" in high standards contexts, the Moore paradox.
- ⁹ Following suggestions of Kripke, Fumerton (1987), and Vogel (1999), Cohen (1999, pp. 72–73) argues that DeRose's Rule of Sensitivity may allow some abominable conjunctions that do not turn on raising epistemic standards.
- ¹⁰ Cf. Schiffer (1996, pp. 326–328) and Kornblith (2000, pp. 28–29).
- ¹¹ Hawthorne (2000, p. 118) correctly observes that there is an "inflection of practical advice" in which we do say such things as "You know you aren't going to win." The inflection signals that the speaker is engaging in overstatement. Cohen (2000b, pp. 135–136) suggests that you "say something true" because the inflection raises the standards. But I do not think we interpret such statements as literally true: we treat them as deliberate overstatements (§VII).
- ¹² Cf. Feldman (1999, p. 107). Contrast Cohen (1987, p. 3; 1999, p. 57), who maintains that the skeptic's claim does not contradict Moore's.
- ¹³ Sosa (2000) assumed that "know" was indexical, and used the invalidity of arguments like the first two to cast doubt on the third. I believe that the intuitive difference between the inferences is evidence that "know" is not indexical.
- ¹⁴ Gerhard Ernst observed in his commentary on this paper that the indexical "on the right" behaves more like "know" in one respect. Indeed, *The tower is on the right in one context but not others* has an interpretation on which it is true. But even this indexical has different linguistic properties, since it creates an ambiguity "know" does not. If we fix the referent of "the right" rather than letting it float, the italicized sentence implies that the tower moves around rather than us, and is then false.
- ¹⁵ It might be thought that "knows for sure" is analogous to "perfectly flat." But we never say things like "Bill knows that he missed but he does not know for sure that he missed" (or worse, reverse the conjuncts). I believe "*S* knows for sure that *p*" is related in meaning to "I will be there, for sure," and expresses the idea that it is certain that *S* knows. Thus when Smith says he knows the flight will land in Chicago, and Mary asks "Do you know for sure?" she is not asking whether he

has some specially strict kind of knowledge, but rather whether Smith is certain that he knows.

¹⁶ Cf. Feldman (1999, esp. 104–107, 111), Klein (2000, p. 113), Kornblith (2000, p. 29), Williams (2000, p. 83).

¹⁷ It is false not because the chemical did not kill all the grass in the universe, but all the grass in the contextually indicated area, namely, our lawn. This example involves the loose use of the quantifier in its restricted, indexical sense. The quantifier can also be used loosely in its unrestricted sense. See Neale (1990, §3.7) for background on the distinction between restricted and unrestricted uses of quantifiers and definite descriptions.

¹⁸ Grice (1975, 1978, 1981), Searle (1975), Sadock (1978), Leech (1983), Levinson (1983), Horn (1992), Neale (1992), and Davis (1998).

¹⁹ Compare and contrast Unger (1984, pp. 6–11, 35–40); Rosenberg (2000, p. 61), and Lehrer (2000).

²⁰ Contrast Cohen (1999, p. 60; 2000b, p. 138), who evidently followed Grice (1975, p. 39) in thinking that what is implicated could never be something entailed by what is said. I observed otherwise in Davis (1998, p. 6).

²¹ Cf. Cohen (1999, p. 65, 80, 83). Compare and contrast Unger (1984, pp. 37–38).

²² For a much more complete answer to this sort of question, see Davis (2003).

REFERENCES

- Annis, D. B.: 1978, 'A Contextualist Theory of Epistemic Justification', *American Philosophical Quarterly* **15**, 213–219.
- Austin, J. L.: 1961, 'Other Minds', in *Philosophical Papers*, Oxford University Press, Oxford, 76–116.
- Cohen, S.: 1986, 'Knowledge and Context', *Journal of Philosophy* **83**, 574–583.
- Cohen, S.: 1987, 'Knowledge, Context, and Social Standards', *Synthese* **73**, 3–26.
- Cohen, S.: 1988, 'How to Be a Fallibilist', *Philosophical Perspectives* **2**, 581–605.
- Cohen, S.: 1999, 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives* **13**, 57–89.
- Cohen, S.: 2000a, 'Contextualism and Skepticism', *Philosophical Issues* **10**, 94–107.
- Cohen, S.: 2000b, 'Replies', *Philosophical Issues* **10**, 132–139.
- Davis, W. A.: 1998, *Implicature: Intention, Convention, and Principle in the Failure of Gricean Theory*, Cambridge University Press, Cambridge.
- Davis, W. A.: 2003, *Meaning, Expression, and Thought*, Cambridge University Press, New York.
- DeRose, K.: 1995, 'Solving the Skeptical Problem', *Philosophical Review* **104**, 1–52.
- Dretske, F.: 1970, 'Epistemic Operators', *Journal of Philosophy* **67**, 1007–1023.
- Dretske, F.: 1981, 'The Pragmatic Dimension of Knowledge' *Philosophical Studies* **40**, 363–378.
- Feldman, F.: 1999, 'Contextualism and Skepticism', *Philosophical Perspectives* **13**, 91–114.
- Fogelin, R. J.: 2000a, 'Contextualism and Externalism: Trading in one Form of Skepticism for Another', *Philosophical Issues* **10**, 43–57.
- Fogelin, R. J.: 2000b, 'Replies', *Philosophical Issues* **10**, 86–93.

- Fumerton, R.: 1987, 'Nozick's Epistemology', in S. Luper-Foy (ed.), *The Possibility of Knowledge: Nozick and His Critics*, Rowman and Littlefield, Totowa, NJ. 163–181.
- Grice, H. P.: 1975, 'Logic and Conversation', in P. Cole and J. Morgan (eds), *Syntax and Semantics*, Vol. 3, Academic Press, New York.
- Grice, H. P.: 1978, 'Further Notes on Logic and Conversation', in P. Cole (ed.), *Syntax and Semantics*, Vol. 9, Academic Press, New York, 113–128.
- Grice, H. P.: 1981, 'Presupposition and Conversational Implicature', in P. Cole (ed.), *Radical Pragmatics*, Academic Press, New York, 183–198.
- Hawthorne, J.: 2000, 'Reply to Cohen', *Philosophical Issues* **10**, 117–120.
- Heller, M.: 1999, 'Contextualism and Anti-luck Epistemology', *Philosophical Perspectives* **13**, 115–129.
- Horn, L. R.: 1992, 'Pragmatics, Implicature and Presupposition', in W. Bright (ed.), *International Encyclopedia of Linguistics*, Vol. 2, Oxford University Press, New York, 260–266.
- Klein, P.: 2000, 'Contextualism and Academic Skepticism', *Philosophical Issues* **10**, 108–116.
- Kornblith, H.: 2000, 'The Contextualist Evasion of Epistemology', *Philosophical Issues* **10**, 24–32.
- Leech, G.: 1983, *Principles of Pragmatics*, Longmans, London.
- Lehrer, K.: 2000, 'Sensitivity, Indiscernibility, and Knowledge', *Philosophical Issues* **10**, 33–37.
- Levinson, S. C.: 1983, *Pragmatics*, Cambridge University Press, Cambridge.
- Lewis, D.: 1979, 'Scorekeeping in a Language Game', in R. Bäuerle et al. (eds), *Semantics from Different Points of View*, Springer, Berlin, 172–187.
- Lewis, D.: 1996, 'Elusive Knowledge', *Australasian Journal of Philosophy* **74**, 549–567.
- Neale, S.: 1990, 'Descriptive Pronouns and Donkey Anaphora', *Journal of Philosophy* **87**, 113–150.
- Neale, S.: 1992, 'Paul Grice and the Philosophy of Language', *Linguistics and Philosophy* **15**, 509–559.
- Nozick, R.: 1981, *Philosophical Explanations*, Harvard University Press, Cambridge, MA.
- Rosenberg, J. F.: 2000, 'Scrutinizing a Trade', *Philosophical Issues* **10**, 58–66.
- Sadock, J. M.: 1978, 'On Testing for Conversational Implicature', in P. Cole (ed.), *Syntax and Semantics*, Vol. 9, Academic Press, New York, 281–297.
- Schiffer, S.: 1996, 'Contextualist Solutions to Scepticism', *Proceedings of the Aristotelian Society* **96**, 317–333.
- Searle, J.: 1975, 'Indirect Speech Acts', in P. Cole and J. L. Morgan (eds), *Syntax and Semantics*, Vol. 3, Academic Press, New York, 59–82.
- Sosa, E.: 1988, 'Knowledge in Context, Skepticism in Doubt: The Virtue of our Faculties', *Philosophical Perspectives* **2**, 139–155.
- Sosa, E.: 2000, 'Skepticism and Contextualism', *Philosophical Issues* **10**, 1–18.
- Unger, P.: 1975, *Ignorance: A Case for Skepticism*, Clarendon Press, Oxford.
- Unger, P.: 1984, *Philosophical Relativity*, University of Minnesota Press, Minneapolis.
- Valdés-Villanueva, L. M.: 2000, 'Contextualism and Levels of Scrutiny', *Philosophical Issues* **10**, 72–79.

- Vogel, J.: 1999, 'The New Relevant Alternatives Theory', *Philosophical Perspectives* **13**, 155–180.
- Williams, M. J.: 2000, 'Is Contextualism Statable?', *Philosophical Issues* **10**, 80–85.

Department of Philosophy
Georgetown University Washington
New North
Washington, DC 20057
USA
E-mail: davisw@georgetown.edu

GERHARD ERNST

IN DEFENSE OF INDEXICALISM:
COMMENTS ON DAVIS

ABSTRACT. Wayne Davis (2004) argues against the thesis that knowledge claims are indexical, and he presents an alternative account of the contextual variability of our use of “*S* knows *p*.” In this commentary I focus on the following three points. First, I want to supplement Davis’s considerations about the inability of indexicalism to deal with “skeptical paradoxes” by considering what the consequence would be if the indexicalist’s explanation of these paradoxes were satisfactory. Second, I am going to take a brief look at Davis’s alternative theory. Third, in the main part of my commentary I try to show that indexicalism may be true in spite of the linguistic evidence Davis presents against it.

1. INTRODUCTION

In his paper “Are Knowledge Claims Indexical?” Wayne Davis poses a serious challenge for the core variety of contextualism, namely indexicalism. The overall structure of his paper is this: Davis starts by identifying what he sees as the main problem contextualism is supposed to solve. He goes on to show that the contextualist solution of this problem is unsatisfactory even if we grant that “*S* knows *p*” indeed behaves like an indexical expression. But, and this is perhaps the most challenging part of the paper, there is abundant linguistic evidence to show that the concept of knowledge behaves quite differently from indexical expressions. So indexicalism not only fails to solve the problem it was designed to solve, it is also in itself an implausible theory. Nevertheless, Davis accepts that there is some truth in contextualism, and he proposes to capture its insights with a theory about what we *imply* when we say “*S* knows *p*.”

Within the bounds of this commentary I can’t go into all the interesting arguments contained in Davis’s paper. So I am going to focus on the following three points. First, since I agree with virtually everything Davis says about the inability of indexicalism to deal with what is called “skeptical paradoxes” in his paper, I only want to supplement his arguments by considering what the consequence would be if the indexicalist’s explanation of these paradoxes were satisfactory. Second, I am not going to discuss the sketch of the



alternative theory Davis offers in any great detail, but I want at least to indicate some preliminary points. Third, in the main part of my commentary I want to consider the question whether the linguistic evidence Davis presents really is fatal for indexicalism. I try to explain why the concept of knowledge might be indexical in spite of what Davis says.

2. THE SKEPTICAL PARADOX

What Davis refers to as “the skeptical paradox” in his paper consists in the difficulty we have in deciding which of the following three sentences we should accept:

- A1 I know that I have a hand.
- B1 If I know that I have a hand, then I know I am not a brain in a vat.
- C1 I do not know that I am not a brain in a vat.

Like Lord Rostov in *War and Peace* who always likes the speech he has just heard best, we seem to change our mind as we go along these sentences. When we hear A1 we want to accept it, when we hear B1 we also want to accept it, and still when we hear C1 we want to accept this sentence as well, even though we know that it is incompatible with A1 and B1. But when we reconsider which of the sentences A1 or B1 we want to give up we again would rather like to accept them, and so on. Davis argues that the indexicalist can’t account for the vacillating attitude we have towards these sentences by invoking the (alleged) fact that the truth conditions of “*S* knows *p*” vary with the contextually indicated standard of evidence, and I do agree with him. Especially I think that he is right in claiming that what is really at issue here (what the skeptical paradox is all about) is a normative question: which standard, the standard we ordinarily accept or the skeptical standard is the *proper* one. Although I think that talking about standards in this context is always a bit vague we could say that what is important here is the assessment of skeptical arguments: we try to figure out whether invoking skeptical scenarios is really enough to provide good reasons to doubt our everyday knowledge claims. That’s at least what philosophers like Descartes tried to do. He did accept that something like the demon hypothesis threatens our ordinary knowledge claims but he thought that we have good reasons to rule such alternatives out. Of course, most of us won’t accept Descartes’ answer. But his question seems to be still with us (even though we are not concerned with absolute certainty any more). Just like Davis I think that there exists a skeptical challenge on

the normative level the indexicalist (at least the indexicalist Davis has in mind) does not even address. As Davis puts it:

When the skeptic questions whether we really know the things we think we know, suggesting that the evidence of our senses is not good enough, it is not adequate to say that we have a different standard of evidence in mind. The question is whether the standard we have in mind is correct. (Davis, 2004, p. 270f.)

One might think¹ that *if* the indexicalist could explain our vacillating attitude towards the inconsistent epistemic triad as he wants to (which, Davis argues convincingly, he can't) this at least would by itself constitute strong evidence for his thesis that the truth conditions of "*S* knows *p*" do indeed vary with the contextually indicated standard of evidence.² But even that is not true. For we might provide alternative explanations for this phenomenon by way of "error theories." Let me sketch briefly how this might be done from the point of view of someone who believes that the standard of evidence relevant for the question whether or not "*S* knows *p*" is true is the skeptical standard. Consider the following three sentences.

- A2 I know that my son Tom is a good boy.
- B2 If I know that Tom is a good boy, I know that he is not a second Jack the Ripper.
- C2 I do not know that my son is not a second Jack the Ripper (because I know that many reliable witnesses claim to have seen him murdering his victims).

Concerning these three sentences Tom's father might very well have a vacillating attitude. The father has always believed that Tom is a good boy. Nevertheless the evidence that he is not is overwhelming. But still the father finds himself again and again questioning this evidence. Even reliable witnesses may sometimes be wrong, and so the father might say: "In this case they *must* be wrong. I don't see at the moment how they can be, but I'll find out because I *know* that my son is innocent." Yet, when considering the evidence again he might be tempted to agree that he doesn't know that his son is not a second Jack the Ripper, and so on. The reason for his vacillating attitude is very easy to see: he simply doesn't *want* to give up his belief that Tom is a good boy. Besides, it is *not so easy* for him to give up this belief even if he wants to, because he is so used to thinking that his son is a good boy, and it is difficult to overcome deep-rooted convictions.

What happens in the case of our skeptical paradox might be very similar: We simply don't *want* to accept that we know almost nothing.

It's like being thrown out of (the epistemological) paradise. And it is *not easy* to give up the belief that I know that I have hands even if I want to. It's such a deep-rooted belief, and every non-philosopher will laugh out loud if I tell him about my ignorance concerning the existence of my hands instead of admiring my intellectual integrity. And yet, how could we exclude the possibility that we are brains in vats? So, just like Tom's father (with respect to (C2)), we try again and again to find a way to deny (C1) even though, in a calm moment, we are ready to admit that it is undeniable. And just like him we find ourselves vacillating between denying (A1/2) and accepting (C1/2) (in a calm moment) on the one hand and accepting (A1/2) and denying (C1/2) (when we are driven by wishful thinking and cognitive habits) on the other.

This would be a psychological rather than a normative explanation for our vacillating attitude towards the three sentences with which we began, i.e. an explanation in terms of what we do rather than in terms of what we are entitled to do. But why should we assume that we are *entitled* to have a vacillating attitude? Making such an assumption would be, of course, begging the question against the non-indexicalist – in this case the non-indexicalist skeptic. Thus, in order to have an argument for the thesis that the truth conditions of “*S* knows *p*” vary with the contextually indicated standard of evidence, the indexicalist not only has to provide an explanation for our vacillating attitude: he also has to convince us that his explanation is better than any alternative explanation.

3. DAVIS'S ALTERNATIVE THEORY

It's the Smith example which provides according to Davis the “solid kernel of truth in the contextualist theory” (Davis, 2004, p. 258). Smith would say that he knows that the flight from LA to New York stops in Chicago when asked by some passers-by at the airport, but when questioned by Mary for whom it is of vital importance to get to Chicago he withdraws his knowledge claim. What needs explanation here is that:

[...] Smith is unlikely to have felt embarrassed about what he said earlier, or to have regretted saying it. He probably did not have the wrenching feeling that he knows less than he thought he did, as we do when we are influenced by skeptical hypotheses. (Davis, 2004, pp. 271–272)

Since Davis does not accept the indexical theory, he tries to account for this example with a pragmatic theory about what is implied by our everyday knowledge claims. According to this theory:

Smith was interpreting “I do” in the first context *loosely*, to mean or imply that *he is sufficiently close to knowing for the purposes at hand*. When the context required greater precision, he stopped speaking loosely. (Davis, 2004, p. 272)

I feel a bit uneasy about this proposal for the following reason. I wonder whether according to this theory we ever claim to know something not loosely. Almost any knowledge claim could be rendered “loose” by imposing stricter standards. But if we are almost always using an expression “loosely” I don’t know what being used loosely really means here. At least *prima facie*, I should think that always using an expression loosely changes its meaning. Nevertheless, I admit that this criticism might be premature since Davis only presents a sketch of his view here. In order to assess his proposal we would have to hear more about the details.

What can be said at this point though is that his alternative to the indexicalist approach is not a very promising line to pursue from the point of view of someone who is interested in analyzing knowledge. Davis’s theory is, of course, not meant to provide an informative answer to the question: When is the sentence “*S* knows *p*” true? But since he wants to account for the variability of our use of this sentence due to changes in the speaker’s context by invoking *pragmatics*, he has to defend an invariantist account of knowledge or at the most a relational form of contextualism.³ And I think that the discussion about the analysis of knowledge during the last forty years has shown that it is very difficult indeed to provide a satisfactory analysis of the concept of knowledge on this basis.⁴ There is at least the hope that an indexicalist form of contextualism (or something very similar) may fare better in this enterprise. So it is well worth taking a closer look at the linguistic evidence which seems to show that indexicalism is doomed to failure.

4. ASSESSING THE LINGUISTIC EVIDENCE

In his paper Davis gives several different arguments against indexicalism which are based on linguistic evidence. I can’t discuss each of them in detail, but I would at least like to explain why I think that the most important of them are not conclusive.

1. Let me start with an argument the indexicalist can handle without much difficulty. It goes as follows. According to Davis the following two sentences might look like stylistic variants of each other:⁵

- a. The sentence “*S* knows *p*” is true in one context, but not in others.
- b. *S* knows *p* in one context, but not others.

Now Davis points out that if “*S* knows *p*” were indexical these two sentences should be as different as the following two which are obviously not even close to being equivalent:

- a. (The sentence) “He is my son” is true in one context, but not in others.
- b. He is my son in one context, but not in others.

Since this is not the case, “*S* knows *p*” doesn’t seem to be indexical. I think the argument is not convincing because not all indexicals behave like the one in the example Davis offers. Consider for instance the following two sentences:

- a. (The sentence) “The tower is on the right” is true in one context, but not in others.
- b. The tower is on the right in one context (namely when seen from one point of view), but not in others (namely when seen from the opposite point of view).

I think in this case (a) and (b) say more or less the same even though they contain indexicals. They might very well be used as stylistic variants of each other. Of course it remains to be seen whether “*S* knows *p*” behaves like “The tower is on the right” in other relevant respects as well.⁶

2. The second argument I want to consider is a much stronger one, and I want to discuss it in some detail. Here Davis points out that there is a difference between “*S* knows *p*” and indexical expressions in the following sentences:

- (1) *S* spoke truly when saying “I am not a man.” Therefore, I am not a man.
- (2) *S* spoke truly when saying “Wayne Davis is a foreigner.” Therefore, Wayne Davis is a foreigner.
- (3) *S* spoke truly when saying “Smith knows that *p*.” Therefore, Smith knows that *p*.

(For expository reasons I slightly changed inference (3). Davis uses “Some people” instead of “Smith,” and “they have a hand” for *p*, but this is of no importance for the present argument.) While the first two inferences are invalid, the last one at least seems to be valid. This clearly poses a severe challenge for indexicalism. Nevertheless, I want to claim that the indexicalist can answer this challenge.

The first step in my argument is to notice a difference between the first and the second inference. While (1) does not even look valid, I think that at least *prima facie* you might be tempted to take (2) as a valid inference. Just consider the following sentence:

- (4) *S* was right when he said “Wayne Davis is a foreigner,” but actually Wayne Davis is not a foreigner.

I think this at least looks like a paradox. Of course we can quite easily remove the air of paradox by pointing to the fact that the context in which *S* takes Davis to be a foreigner might be different from the context in which (4) is uttered. John might have used (4) in America, while *S* said that Wayne Davis is a foreigner during a visit of Davis in Germany. The first part of an explanation why (4) looks valid is that it does not contain any such explanation. And the same holds for (2). Inference (1) on the other hand does contain the relevant hints. It is explicitly mentioned who makes the utterance “I am not a man,” namely *S*, and on the natural assumption that almost nobody talks about himself in the third person we infer immediately that the person who uses (1) usually will differ from *S*. Therefore it is obvious that two different points of view might be involved, and so (1) does not even look valid.

The second part of the explanation why (4) looks like a paradox and (2) looks valid is that (special situations aside)⁷ we normally use indexicals without further explanation only when the person we address or talk about shares our context. (The indexical “I” is rather special in this respect.) For example we normally say things like “The tower is on the right” when the person we address is looking in the same direction as we are. And we would not without further explanation say things like “John spoke truly when saying ‘This place is hell’” if we would not share John’s context. Therefore it is natural to assume that the person using (4) and (2) shares the same context with *S*. And on this assumption (4) *is* a paradox and (2) *is* a valid inference.

I think that these considerations show that the fact that we can’t see the invalidity of (3) immediately without further explanation does not yield a conclusive argument against the indexicality of “*S* knows *p*.” If we normally use the concept of knowledge only when we share a common point of view, it should not come as a surprise that (3) looks valid. The sentence “*S* spoke truly when saying ‘Smith knows that *p*,’ but it is not true that Smith knows that *p*.” looks like a paradox, but so does the sentence “*S* spoke truly when saying ‘The tower is on the right,’ but it is not true that the tower is on the right.”

Further explanation can remove the air of paradox in the latter case. Is it possible to remove it in the former?

In order to show that it is, I want to return to the example Davis himself accepts as providing evidence for the variability of our use of “know”: the Smith example. Smith, having just looked at his flight itinerary, says “I know that the flight from Los Angeles to New York stops in Chicago” when asked by some passers-by. And he says “I don’t know whether the flight from Los Angeles to New York stops in Chicago” later when Mary makes it clear to him that it is of vital importance for her to get to Chicago. Still, Davis agrees, he need not feel embarrassed about what he said earlier. To strengthen this point let’s assume in addition that later still Smith is informed about the fact that the flight did indeed stop in Chicago. It is now perfectly clear that he will not feel embarrassed about what he said earlier. But I don’t think that he will feel embarrassed about what he said later either. After all, it was important for Mary to get to Chicago, and he was not in the position to rule out the possibility that his flight itinerary was wrong.

Now imagine, when informing Smith about the fact that the flight did stop in Chicago, we ask him: “Did you know that the flight would stop in Chicago or didn’t you?” In this situation I think it would be perfectly reasonable for him to answer something like this: “Well, you could say I did and you could say I didn’t: considered from one point of view (i.e. if we assume a common or garden standard to be relevant) I did know that the flight would stop in Chicago, but considered from another (i.e. if we assume Mary’s standard to be relevant) I did not know that the flight would stop.” In this situation the question whether he *really* knew that the flight would stop in Chicago seems to be very similar to the question whether the tower really is on the right. He really knew it given low standards and he really didn’t given high standards. There seems to be no fact of the matter about knowing here (only about knowing from a certain point of view) as there is no fact of the matter about being on the right (only about being on the right from a certain point of view). Simply to assume at this point (contrary to appearances) that there *must* be a fact of the matter about knowing in this case would be, of course, simply to assume that indexicalism is false.⁸

Inference (2) is invalid because there is no fact of the matter about being a foreigner (only about being a foreigner from a certain point of view). And, as the Smith example shows, (3) can be invalid for a similar reason. From the fact that Smith spoke truly when saying “I

know that the flight from Los Angeles to New York stops in Chicago” (i.e. from the fact that Smith knew that the flight from Los Angeles to New York stops in Chicago from one point of view, i.e. if we assume a common or garden standard to be relevant) it doesn’t follow that Smith knew whether the flight from Los Angeles to New York stops in Chicago (from another point of view, i.e. if we assume Mary’s standard to be relevant).⁹

Given this explanation I think we can see that (3) is indeed in general an invalid inference. Of course, much more would have to be said about why it *looks* valid and what a *point of view* exactly is in this context, but in order to be able to do that I would have to describe which theory of knowledge I want to defend, which I can’t do here.¹⁰ When we proceed from inference (1) to inference (3) the indexicality becomes more and more difficult to detect. But that does not preclude it from being there.

3. Given these considerations we can immediately deal with another of Davis’s arguments: In the case of knowledge it seems to be very easy to raise standards but very difficult to lower them again. For example we only have to mention skeptical hypotheses, and it seems to be impossible to return to our ordinary standard. There’s no analogy to that in the case of other indexicals.

I think the problem about this argument is that raising the standards to the level of the skeptic is untypical for the phenomenon of changing standards in general. Consider again the Smith example. As soon as he knows that the flight did indeed stop in Chicago he can raise the standard (i.e. assume Mary’s standard to be relevant) or lower the standard (i.e. assume a common standard to be relevant) as he likes. Focusing on the possibility of error raises the standard, focusing on the fact that in this case no error occurred lowers it. He can adopt both points of view. What is special about mentioning the skeptical hypotheses is that we are hard put to find arguments to rule them out. But imagine Descartes’ or Putnam’s or some other reply to the skeptic would be convincing. If this were the case again we could adopt two different points of view concerning the knowledge of ordinary people who never thought about skeptical scenarios in the first place. We could say: They don’t know how to rule out skeptical hypotheses, therefore they don’t know that they have hands (higher standard). But we could equally say: Skeptical scenarios are ruled out (which is shown by our philosophical arguments), therefore you don’t have to rule them out in order to know that you have hands (lower standard). There is no “downward stickiness of epistemic standards” as soon as we have all the relevant information. Of course, as long as

we are engaged in an actual inquiry (which many think we are concerning skeptical hypotheses) we can't say that we know a lot given our ordinary standards, because we don't know whether our ordinary beliefs are really true!

4. The last argument of Davis I want to discuss concerns the question whether the indexicalist can account for the fact that there really is a skeptical challenge. As I said: I don't believe that indexicalism alone can answer that challenge. But according to Davis the indexicalist cannot even account for the fact that there is a challenge, because it seems as if he has to interpret the thesis "No one knows anything" as obviously true or false. Either it expresses an obvious falsehood (in an everyday context) or an obvious truth (in a skeptical context). But this is not true. Indexicalists need not accept that "No one knows anything" is obviously true in a skeptical context. At least such philosophers as Descartes tried to show that this sentence is false in a skeptical context. And it is hard to believe that he tried to refuse an obvious truth.

Let me summarize: (1) I think that Davis is right in claiming that indexicalism alone is not very effective in dealing with the skeptical challenge although I believe the indexicalist need not deny that there is a challenge. (2) Davis's pragmatic theory is not a promising alternative to indexicalism for someone who is interested in analyzing knowledge. (3) I hope to have shown that the major arguments Davis provides against indexicalism are not conclusive. Since in my opinion indexicalism (or something very similar) alone has the potential to provide at least part of an adequate analysis of knowledge it would be well worth discussing the arguments in much more detail than I could in this commentary. I think contextualists should change their focus: from dealing with skepticism to analyzing our concept of knowledge.

ACKNOWLEDGMENTS

For very helpful comments on an earlier draft of this paper I would like to thank Erich Ammereller, Elke Brendel, Wayne Davis, Christoph Jäger, Christian Seidel, Sophie Wiesner, Mark Textor, and an anonymous referee. I also want to thank Elke Brendel and Christoph Jäger for organizing the conference on "Contextualist Approaches to Epistemology" (University of Mainz, September 4–6, 2003) at which this commentary was presented.

NOTES

- ¹ This thought is the point of departure for Davis in his paper.
- ² The argument, of course, being an inference to the best explanation.
- ³ Relational contextualists hold that whether “*S* knows *p*” is true depends on the context of the subject *S* (which may be very different from the speaker’s context).
- ⁴ In Ernst (2002) I argue that it is actually impossible to provide an analysis of knowledge on that basis.
- ⁵ At this point Davis refers to Cohen who sometimes even uses them as stylistic variants. But note that contextualists need not and usually would not accept them to be really equivalent. Only the first of them represents their view correctly!
- ⁶ Referring to this passage of my commentary, Davis points out in (2004, note 14) that “The tower is on the right in one context but not others” creates an ambiguity “know” does not. In fact, I think, “know” does create a similar ambiguity although this is less obvious than in the case of “on the right.” I will come back to this point. Cf. note 9.
- ⁷ For example, I take telephone-calls to be a special form of communication.
- ⁸ Davis, of course, provides an alternative explanation for the Smith example. But only if indexicalism has been shown to be false already does the need for an alternative explanation arise!
- ⁹ If this explanation is sound “*S* knows *p* in one context, but not others” creates the same ambiguity as *The tower is on the right in one context, but not others* does (cf. note 6). Davis says (in his note 14): “If we fix the referent of ‘the right’ rather than letting it float, the italicized sentence implies that the tower moves around rather than us, and is then false.” Fixing the referent of “the right” amounts to fixing a point of view: If we fix a point of view, the italicized sentence implies that the tower moves around which is not what we wanted to say when we said “‘The tower is on the right’ is true in one context, but not in others.” But now we can see that something similar is true in the case of knowledge: If we fix the point of view considered from which we ascribe knowledge, the sentence “*S* knows *p* in one context, but not others” implies that whether or not *S* knows *p* depends on contextual factors other than our point of view which might be true, but which again is not what the indexicalist wants to say when he says: “The sentence ‘*S* knows *p*’ is true in one context, but not in others.”
- ¹⁰ Cf. Ernst (2002).

REFERENCES

- Davis, W.: 2004, ‘Are Knowledge Claims Indexical?’, *Erkenntnis* **61**, 257–280.
 Ernst, G.: 2002, *Das Problem des Wissens*, Mentis, Paderborn.

Seminar für Philosophie
 Logik und Wissenschaftstheorie
 Ludwig-Maximilians-Universität München
 Ludwigstraße 31
 80539 München
 Germany
 E-mail: gerhard.ernst@lrz.uni-muenchen.de

VERENA GOTTSCHLING

KEEPING THE CONVERSATIONAL SCORE: CONSTRAINTS FOR AN OPTIMAL CONTEXTUALIST ANSWER?

ABSTRACT. Conversational contextualism states that the truth-conditions expressed by knowledge-attributing sentences vary relative to the context of utterance. This context is determined partly by different standards the person involved must meet in order to make the sentence true. I am concerned with the question of how these standards can be raised or lowered, and especially what happens to the standards and the conversational score when parties in a discussion push the conversational scores in different directions. None of the available options for an answer seems satisfying. I argue that this results from a misunderstanding of the characteristics of the situation at hand.

1. INTRODUCTION

Conversational contextualism states that the truth-conditions of a knowledge ascription are relative to a context and that this context determines certain standards. I am here concerned with theories according to which the truth-conditions of knowledge ascribing sentences are sensitive to certain facts about the conversational participants in that context. For each context and knowledge ascription, there is an epistemic standard, how strong one's epistemic position has to be with respect to the proposition P in order to count as knowledge that P . The truth-value for a given knowledge-ascribing sentence can vary because its content can vary. These standards can be raised or lowered during the conversational process. The specific question I address is this: (Q) What happens in situations where the standards change during the conversation, but only one participant accepts this change? The other participant refuses the change of standards, i.e. both parties in a discussion push the conversational scores in different directions. In a recent paper addressing this issue, Keith DeRose (DeRose, 2004) develops and defends the "GAP" view as the best answer for a conversational contextualist. I begin by describing the problem, the available solutions, and the arguments DeRose presents for the GAP view.



After drawing attention to problematic aspects I see in his proposal, I argue that the real problem is caused by a faulty presupposition and that we will not be able get an adequate answer to *Q* until we correct this presupposition.

2. THE PROBLEM

To simplify matters, let us assume we have only two conversational partners, a skeptic *S* and a commonsensian *C*. They are accompanying their friend *R* to London Heathrow airport and discussing whether *R* knows that the 9:00 a.m. plane to Chicago is leaving on time (proposition *P*). *S* makes a statement that raises the standard to a special degree: “*R* called the airline 30 minutes ago and asked them about it. But *R* cannot rule out the possibility that they have changed the schedule since then. Thus, *R* does not know. We better take a look at the departure board.” As a result, given the standards put in place by *S*’s remarks, *R* would not count as knowing *P*. But *C* refuses to raise the standards. His reaction is “Come on, they wouldn’t change the schedule in just thirty min. They guaranteed *R* the plane will leave on time. Thus, *R* does know.” and *C* persists in his refusal to raise the standards. This is the beginning of a fierce discussion as to whether *R* knows that *P*:¹

(a) *C*: “*R* knows that *P*.”

(b) *S*: “*R* does not know that *P*.”

What is happening in cases like this, where the content of a context-sensitive term like “know” is used divergently in a personally indicated way? How should we describe these kinds of conversations? What happens to the truth-conditions of the knowledge attributing claims in these kinds of situations? The problem refers to an overall fundamental problem: in case of divergent personal standards, how is the shared conversational score determined? The challenging task regarding these questions is to determine, what a contextualist *should* say. Are the higher standards relevant for the scoreboard, is it usually enough to keep the old standards, or do both parties have and maintain their own standards and their own scoreboard? The problem arises because there are two main intuitions here: (IP) On the one hand, we want to say that both participants assume that their own standards are the correct standards and that they meet those standards; therefore, their claims (a) and (b), respectively, are both true. In addition, the conversational participants maintain their standards

during the conversation. Let us call this the *intuition of persisting individual standards*. This intuition is the basis for our belief that both participants somehow talk at cross-purposes. (IC) On the other hand, we want to say that both claims contradict one another. After all, that is the topic of the whole discussion. The sentences (a) and (b) seem contradictory. Let us label this the *intuition of contradiction*. This intuition is also deeply entrenched. The core problem is that these two intuitions seem to be incompatible.

Let us systematically assess our options (see Figure 1). Option (1): The idea of different scoreboards is tempting. Each speaker has a scoreboard “in the head”. The score of the whole conversation is a function of the different scoreboards involved. In this view, both conversational partners – speaker and hearer – are correct. In these cases there is no common epistemic standard, thus both are right regarding their own epistemic standards and scoreboards. Because both participants pull the conversational score in different directions, there is no shared standard. A second option is that there is a single scoreboard. In that case we have to commit ourselves to a position with reference to how the score is set. We have three subversions of this view: (2i) the skeptic sets the score, (2ii) the commonsensian sets the score, or (2iii) both partners play a role in determining this scoreboard. Finally, (3) we could state that in these cases there is no scoreboard because the standards of both participants are too divergent.




Figure 1. Overview: The systematic options.

Which contextualist alternative is the best? In the literature, views (1) und (2i) seem to be the chief favorites. In a recent paper DeRose (DeRose, 2004) argued that both favored views in the literature have one thing in common – they entail that the skeptic says something true, i.e. that (b) is true. That makes them somehow “skeptically-friendly,” and DeRose holds this responsible for the unattractiveness of contextualism to many participants in the debate. In addition, he argues that none of the options (1)–(3) is able, so far, to accomodate both intuitions. DeRose argues that another view, (4) the GAP view, gives us a better grip on both problems (DeRose, 2004).

3. FLAWS OF THE STANDARD OPTIONS

Unfortunately, the view with the most intuitive appeal, the “different scoreboards” view, according to DeRose, injures the intuition of contradiction. On the different scoreboards view, each participant of the conversation has his/her own personally indicated content because of the context-dependent term “know” and her own personal scoreboard. As a result, both speak the truth and do not contradict one another because there is no shared scoreboard. But in fact, *S* and *C* both perceive themselves to be contradicting one another. After all, both speakers indicate that they understand and plan their claims and denials to be contradictory. This is the issue at stake in their whole discussion.

Therefore, to do justice to the intuition of contradiction, we must admit that there is at most one common scoreboard. But who sets the score? The problem with both the “skeptics managed” view and the “commonsensian managed” view is that we violate the intuition of persisting individual standards. Because one of the speakers sets the score, the truth-values of the other speaker’s claims do not cohere with her own standards. Conversational contextualists frequently make claims that seem to back up the view that the skeptic sets the score. Lewis argues that when the skeptic brings up a new possibility, there is a new relevant alternative available, which raises the standards. Thus, the high standards of the skeptic prevail, and “The commonsensical epistemologist must concede defeat” (Lewis, 1979, p. 355). It is tempting to interpret Lewis here as saying that in cases like the one we are interested in, the stubborn commonsensian says something false. He uttered something true with respect to the former score, but after the skeptic brings up the new possibility, the score is changed. If we state that the other participant sets the score, we run into the same problem. In “Elusive Knowledge”, David Lewis de-

fends that the standards can be lowered in a conversation with high standards, when something is said that is true under the lower standards, and nobody rejects, the standards are lowered. This seems to imply that when the participant is refusing to change the standard, it does not change. DeRose suggests that a similar position, he introduces as “veto power”, is “tempting” about our debate. Our commonsensian uses his veto power. He does not accede to the raising of standards, therefore, the sceptic is not succeed in changing the standards. Does the commonsensian in our example set the score? It is important to be aware that in Lewis’s example referred to here, the issue at stake is whether the standards go *down*. Thus, the commonsensian tries to lower the standards, the *skeptic* is reluctant and refuses the change. The skeptic’s veto gets in the way of *lowering* the standards. Thus, the skeptic again sets the standards for the conversation. Veto power to impede the raising of standards is much more problematic. But the question is not about the mechanisms of changing the standards, the question is how we should estimate the resulting situation. In fact, the label “veto power” seems misleading here because, by this classification, we do not systematically exhaust all the options we have. Furthermore, in this way we mix up two different classifications: First, there is a “level classification” in our conversation, skeptic versus commonsensian; second, there is a “classification of the functional roles”, the one willing to change the standard versus the reluctant opponent using veto power. To increase clarification and to register all systematical options, I classified the options as “skeptics managed” view versus “commonsensian managed” view. The final subversion of the “one scoreboard” view is the “balanced scoreboard” view.² According to this view, the score of course changes as the conversation progresses. It moves up when the skeptic speaks and drops when the commonsensian answers. Thus, the maneuvers made by all participants are registered, but at any given time there is only one single score. There are two possibilities here: first, we can assume that this up and down movement happens with every utterance – this view is very similar to the different scoreboards view and, according to DeRose, injures the intuition of contradiction. Second, we assume that after two or three standards-changing maneuvers of non-cooperation, the score reaches a stable state (see Figure 1, 2iii) (DeRose, 2004). This version of the view is more promising but loses one of the two main intuitions as well: the intuition that both partners are right regarding their own standards.

Yet another possible answer is to state that there is no correct shared scoreboard in our example because claims involving the relevant

context-sensitive terms are neither true nor false. Because of the different standards of the participants, the key term “know” has different meanings, and the sentences have different truth-conditions. For that reason “semantic hell breaks loose” as DeRose colorfully puts it. Since there is no common scoreboard, all claims involving the relevant terms from the skeptic and the commonsensian lack truth value. They are neither true nor false. What about the intuitions with which we started? Our condition that both parties say something true regarding their own persisting standards is probably fulfilled. But are the participants contradicting one another? I would answer that that depends on how you understand the intuition of contradiction. Is it enough that both speakers insist on their own standards? Apparently not: We rejected the different scoreboards view because the intuition of contradiction is violated. A contradiction presupposes that there is a shared conversational scoreboard, a controversial proposition but different truth-values. In the “no-scoreboard” view, the intuition of contradiction seems violated for a similar reason. Each participant says something that is neither true nor false, because the meaning of the relevant terms differs. There is no common scoreboard, therefore no contradiction.

To summarize, all standard views lose one of the main intuitions with which we started. All “one scoreboard” views violate the intuition of persisting individual standards (IP). The different scoreboards view and the no scoreboard view in contrast fall short of the intuition of contradiction (IC).

4. DEROSE’S GAP VIEW: A SOLUTION?

4.1. *The GAP View*

DeRose’s preferred answer to the problem is the GAP view: “*R* knows that *P*” is true if both personally indicated standards are met by *R*, it is false if *R* fails to meet either set of standards. In case *R* meets just one standard it is neither true nor false.³

<i>S</i> ’s standards	<i>C</i> ’s standards	Scoreboard
Met	Met	T
Met	Failed to meet	U
Failed to meet	Met	U
Failed to meet	Failed to meet	F

The advantages of the GAP view are obvious. In cases of great divergence, the GAP view's result is a huge gap – most of the claims in such a “gappy” conversation have no truth-values. In effect, we end up with the no scoreboard view, but admit that (1) “*R* know that *P*” is false if *R* does not meet the standards of at least one participant in the conversation, and (2) “*R* know that *P*” is true if *R* meets the standards of both participants. In cases of small divergence, the gap is smaller – therefore, the chances of getting a truth-value are much better. There are two personally individuated standards (in column one and two) but only one scoreboard; both standards determine the truth-values of the claims. DeRose states that thus both intuitions we started with are respected. In addition, the view is less skeptic-friendly.

I am puzzled about the details of this purported solution, its implications and how such a proposal can be successfully carried out. Let me start with some worries.

The GAP view is symmetrical, but there is an asymmetry to be found between raising and lowering standards.⁴ DeRose is concerned with the question of how we should describe the status quo after the non-cooperation in the conversation occurs. In principle, the common conversational score could be constituted differently in cases of lowering versus raising standards. But in classical conversational contextualism, it is impossible to say how, exactly, the standards change in a given conversation. That originates from the fact that some conversational rules, like the rule of attention and the rule of resemblance, depend on facts about the speaker and the hearer,⁵ others do not, like the rule of actuality and the rule of belief. That means, in effect, that the course of the conversation cannot be predicted, even if we know all relevant external factors. So, we cannot expect to get an explanation for the mentioned asymmetry. Is this a problem for the GAP view? I do not think so. On the contrary, it rather supports this position. It shows that changing the standards is complex, resisting a change is easier in one direction than in the other. And in cases of odd conversations like the ones we are interested in, it predicts that conversational mechanisms are violated. For that reason, there is no shared conversational score, and we end up in the gap. In addition, nothing prevents the GAP view in principle from identifying more “fine-grained” rules.

4.2. *Is the GAP View a New Solution?*

As already mentioned, conversational participants typically adjust their standards to each other. If they are not sure what their partner

intended to claim, they ascertain by asking and, as a result, the key terms come to be used in the same way. In our example, however, this is precisely what does not happen. The personally indicated contents of the speakers differ, but both nonetheless take themselves to be contradicting each other. Both participants refuse to adjust to the other's personally indicated content. Thus, we do not have a typical conversation here, but rather a pretty uncommon one.⁶ In fact, the most central term in the conversation, i.e. “knows”, is employed differently. We have, as a result, a divergence. The question *Q* we address concerns cases of disagreement, not cases in which the claims according to both standards are respectively true or false. Why is the “no scoreboard” view committed to claim that *all* ascriptions containing the term “know” are neither true nor false? Regarding the cases of persisting disagreement both views do not differ at all. Regarding the problem we started with, the GAP view is not an alternative but a generalization of the no scoreboard view to cases with small differences as well.

The GAP view's promised improvement is a uniform explanation of cases of small and large differences. Let us take a closer look at these cases. DeRose explains small divergence with an example in which “here” is used by two conversational partners slightly differently. The word “here” when used by one participant designates a slightly larger area than when used by the other participant. DeRose does not give us an example with knowledge ascriptions but seems to assume that doing so is unproblematic.⁷ Transferred to our case, that would mean the criteria for ascribing knowledge that *P* to *R* are almost the same, albeit not exactly. Let us assume that for *S* it is sufficient to check the departure board every 27 min whilst *C* is convinced that checking it every 30 min is enough. Both deny that relying on the information from a phone-call is sufficient for knowing *P*, both rule out that they have to talk to the pilot to acquire knowledge. According to the GAP view, the claim is true if both personally indicated standards are met, it is neither true nor false if it fails to meet only one standard, and false if it fails to meet either set of standards. That only means that the slightly stronger standards of the skeptic must be met for a true knowledge ascription and the slightly lower standards must fail to be met for a false one – thus the respectively “more extreme” standards determine whether the knowledge ascription is true or false. But there is a difference to the other one scoreboard views because the personally indicated standards of the individual positions do not necessarily match the truth-value on the scoreboard: when the personally indicated contents differ, the personally indicated standards of

the participants and the truth-values of their claims do not correlate (because there is no truth-value).

I am troubled by the details: consider again our example. Does that entail that there are almost no cases in which *S* comes into conflict with *C* regarding knowledge-ascribing or denying cases in their conversation? Yes, probably. But we are dealing with a conversation with disagreement. *S* strongly denies *R*'s knowledge that *P*, *C* stubbornly claims that *R* does possess knowledge that *P*. Thus, there is a difference in the personally indicated contents of the individuals, and the truth-values of their claims do *not* match even if the criteria only differ slightly. We always end up in the gap. But how can we decide in any concrete situation whether a claim has a truth-value or not? Here DeRose gives us an illuminating answer. His solution is the no scoreboard view, modified by the, in all other views, uncontroversial cases of the peripheral regions, in which, according to both standards the claims are true or false. We ruled out the other views because they conflicted, by some means or other, with one of the intuitions. Therefore, this solution, in a sense, handles both intuitions – the intuition that something is deeply wrong in cases of large divergence and the intuition that in case the truth-conditions that both have in mind are fulfilled, there is a truth-value. An account that clearly distinguishes between cases of small and large divergence, not only has to draw a clear line between the two, it is more complicated, in that it must give different rules for these different cases. But it is also faced with the task of justifying its disunification. I see problems for both challenges.

4.3. *The Main Intuitions and the GAP View*

What about the intuitions we started with? DeRose argues they are fulfilled. I disagree. In my opinion, a central issue is how we understand the intuition of contradiction and the intuition of persisting individual standards. Regarding the intuition of contradiction, we already pointed out that having different truth-values for *C*'s and *S*'s utterances is not sufficient. In cases of divergence in the personally indicated contents, we need at least a shared scoreboard. The GAP view proposes a shared scoreboard, but no truth-value. Is that enough to justify the intuition of contradiction? I doubt it. A shared scoreboard without entries in the relevant arrays, seems only a marginal improvement to no scoreboards in these cases. DeRose admits that the no scoreboard view is convincing to him in cases of large divergence. As we saw, regardless of how large the divergence

is, the critical cases are those with disagreement; i.e. claims involving the relevant terms from the skeptic and the commonsensian are neither true nor false. Thus, the GAP view seems to get into the same difficulties as the no scoreboard view. If the no scoreboard view is guilty of violating the intuition of contradiction, the GAP view has exactly the same result in the case at hand. Thus, we do not have contradicting claims, i.e. the intuition that both participants are contradicting one another seems to be violated. In case both indicated contents give us the same truth-value, we get the same truth-value for the shared scoreboard. As a result, the truth-value of a claim is never *opposite* to the speaker's personally indicated standard. But, in a sense, there is a deep tension here because whenever there are differently indicated contents there is no truth-value at all.

Regarding the intuition of persisting individual standards, we are no better off. What does it mean to say that both participants maintain their individual standards during the conversation? Is the intuition of persisting individual standards fulfilled if there is no truth-value in the scoreboard? It depends on how you interpret the requirement, that regarding the personal standards the individual claims are true. If you understand it as stating that there is no opposite truth-value in the shared scoreboard – the answer is “yes”. If you understand it as stating that the truth-value in the scoreboard matches the truth-value of the individual claim – the answer is “no”. In my view, the intuition should be understood as saying that the truth-value of the claim is *in consonance* with the individual standard of the person during the whole conversation. Otherwise, it seems hard to say the participant *maintains* her individual standard. According to the GAP view, this maintenance is not guaranteed: it may also be the case that there is no truth-value to the claim, but according to individual standard the claim would be true. Therefore, the intuition is not given up, but considerably weakened.

In case of minimal difference in the standards – both conversational partners demand almost the same conditions must be fulfilled in order to guarantee that the knowledge claims are true – there may be some situations in which a claim may have no truth-value. DeRose admits that this is not the “best job possible”. This has the unwelcome result that the truth-value of my claim depends on something other than my own standard, it depends on the standard of my conversational partner. Imagine my standards and my knowledge ascription “*R* knows that *P*” stay fixed, the only thing that changes are the standards of my conversational partners. It

could happen that my claim is true relative to my standards but has no truth-value in that situation. At least on this view, the truth-value of a claim is never contrary to the speaker's personally indicated standard. It has the same truth-value, or there is no truth-value. Because the gap is small, the prospects of achieving a truth-value are good. This truth-value is determined by the standards of the "more extreme" standard – for true claims the skeptic sets the score, for false claims the commonsensian. Remember, the truth of *S*'s knowledge claim "*R* does not know that *P*" depends on whether *C*'s standards are met by *R*. Let us assume nothing changes, but (without further reason) *C* decides this minute to slightly change her standards without communicating so. As a result, *S*'s knowledge ascription has immediately a different truth-value when uttered a second later. Because *C*'s standards silently change, *S*'s claim expresses a different proposition. *S* would not even be aware of this change. As a result, we cannot decide in any concrete situation whether a claim has a truth-value or not. In fact, we can never be sure that our own claims even have a truth-value, because we cannot know for sure our claim meets the standards of our conversational partner.

It seems to me the advantage of having a uniform solution is accompanied with a serious drawback in that both initial intuitions are considerably weakened. This results from the incompatibility of the two intuitions. A uniform solution that tries to handle both intuitions can only do so by weakening both of them. I doubt whether we should bite that bullet.

Therefore, one disadvantage of the view already pointed out is that the GAP view falls victim to the same problem we stated previously for our other options. At the very least the intuitions are weakened to a great and alarming extent. As I argued, if the different scoreboard view runs into the problem of violating the intuition of contradiction, the GAP view is in the same position. At least, there is no advantage of the GAP view to be stated.

In cases of *agreement* there is a truth-value, which is in consonance with the standards of both participants. But this is predicted by all the remaining views (see Figure 1). But that was not the problem with which we originally started. We wanted to know what happens to the common scoreboard in cases of no agreement and what the scoreboard entry is in cases where the personal standards assign different truth-values to the knowledge ascribing sentences. And in these cases there is no advantage.

5. THE SCOREBOARD PROBLEM

The core problem is this: How is the shared conversational score constituted in case of divergent personal standards? The answer to this problem depends on what you mean by scoreboard and conversational score. I (like DeRose) take the claim that speakers have *personally indicated content* to mean they have certain standards that their conversational maneuvers put in place (or have some tendency to put in place). In addition, according to Lewis, speakers have their *own personal scoreboard*: The truth-conditions of each speaker's use of "knows" is particular to that speaker and presumably matches that speaker's personally indicated content. *Conversational score*, on the other hand, is something quite different: it is a complex function of the speaker's intentions, the listener's expectations, presuppositions of the conversation, salience relations, etc. (Lewis 1979). On my reading, Lewis would describe our example as involving two scoreboards but only one conversational score. Determining the score in a given conversation is complicated, especially in cases where there is no shared meaning of important terms. In contrast, DeRose proposes two standards but one shared scoreboard. But, of course, we are confronted with the same problem, namely, that of describing precisely how the two standards determine the common scoreboard. It seems the discussion of whether we have none, one, or different scoreboards just shifts the basic problem. To formulate it differently: The answer to the question *Q* we started with seems essentially to depend on what we mean by "scoreboard": a shared board for the conversation, like DeRose describes, or something personal each participant has "in his head". In the case of personal individual scoreboards we are bound to a "different scoreboards" view. In my opinion, herein lies the source of the problem raised by Richard Feldman (2004), namely, that it makes a difference for DeRose whether someone silently believes or utters a thought because he has to decide whether the pure thought influences the shared scoreboard. Lewis deals with individual scoreboards; therefore, beliefs and background knowledge play a role. Thus, Lewis's characterization has an advantage at this point. For he clearly distinguishes three things: First, the personally preferred standards to be found in the personally indicated content; second, the individual scoreboard in which the conversational process but also individual states like intentions, expectations, etc. are registered; and third, the conversational score characterized by the aforementioned complex function.

On closer examination, it is not even clear, whether the different scoreboards view and the no scoreboard view in fact differ. According to the different scoreboards view, both participants have their own *personal* scoreboard and, as a result, both speak the truth. According to the no scoreboard view, both participants say something that is neither true nor false because the truth-conditions of the relevant terms differ, and therefore, there is no *common* scoreboard. It seems we are concerned with different questions here. In the first case, we talk about the individual scoreboards “in the head” of the participants. The different scoreboards view cannot even consistently be interpreted as claiming there are multiple *shared* scoreboards. In the case second, we claim there is no shared scoreboard mirroring the whole conversation available, but the no scoreboard view is at least silent about the possibility that both have their own personal scoreboards, which apparently differ. Therefore, the different scoreboards view and the no scoreboard view are not conflicting.

DeRose does not address the question of how, exactly, the individual scoreboards are registered in the common conversational score. His concern lies with the basic parameters for an optimal answer, i.e., that such an answer has to respect both intuitions with which we originally started.

6. BACK AT SQUARE ONE

It seems we are back at square one. The question *Q* with which we began was: How should we best understand conversations where the personally indicated contents of a context-sensitive term like “know” differ? We considered the available options in turn. It seems impossible to do justice to both of the basic intuitions with which we initially began. I want to conclude by proposing, not answers but some signposts pointing in the right direction. Obviously, we have a further option in the current situation: perhaps one of the intuitions is the culprit, and we should give it up? My claim is that these two intuitions are far from being on a par. I introduced the intuition of persisting individual standards (IP) as having three parts: first, both participants assume their standards are the correct standards; second, they meet and maintain their own standards; and third, therefore, their respective claims (a) and (b) are both true. All three claims are compelling. Moreover, IP seems an intuition we cannot give up because it follows from an essential intuition for contextualist accounts. The whole idea that knowledge terms are context sensitive and thus, a

knowledge claim can be true due to one standard, even if it is false according to another stronger standard, is the grounding for the intuition of persisting standards.

Giving up this intuition is clearly not the best strategy for contextualists. What about the intuition of contradiction (IC)? This intuition also seems quite strong. Nevertheless, I find it a more suspect candidate. In the next and final section, I will first explain why embarking on this strategy seems the most promising contextualist response to our worry. Then, I argue that we are misled by interpreting this intuition as being on equal footing with the other one because of a severe misunderstanding of what contextualism claims. If we are successful in undermining IC, the result will be that three views are back in the game: the different scoreboards view, the no scoreboards view and the GAP-view.

7. DEBUNKING THE INTUITION OF CONTRADICTION

I admit the intuition of contradiction is initially quite compelling and difficult to give up. But what exactly is the intuition we do not want to abandon? DeRose is silent about this issue and offers little in defense of the intuition. It would certainly be unedifying were the entire debate reduced to a clash of intuitions. I claim IC needs to be debunked, while conversational contextualists insist the intuition has to be met. One strategy could be to argue that the intuitions contextualists typically take for granted are far from being universal. Nichols et al. (2003) present empirical evidence that the intuitions underlying skeptical and contextualist premises are not as robust as typically assumed by contextualists but vary with socio-economical status, and cultural and educational background. This seems a promising empirical strategy for our intuition of contradiction, as well. Nonetheless, contextualists seem to have various defense strategies against this kind of argument available. They can simply debunk the semantic intuitions of all those who use the term “knowledge” in a loose way. They can also argue that the experiments are constructed in a way that the real intuitions are not “triggered”. My goal is more modest. I aim simply to explore some of the apparent commitments of contextualism for the issue at stake. What precisely is the intuition we are talking about? In my view the exact formulation of the intuition of contradiction is central. The intuition is supposed to be that both partners contradict one another. Obviously, they think they do. And they intend to. But do they? Is the intuition of contradiction

grounded in the fact that in skeptical contexts we somehow have the intuition that we lose knowledge we had before the skeptic brought up several error possibilities? In that case, the intuition is that both claims *in fact* contradict one another. But on a closer reading of contextualism, this is false.

For a contradiction, the knowledge ascriptions of the form “*R* knows that *P*” have to be identical, but the truth-values differ. Contextualists often seem to state that the contents of the knowledge ascriptions remain identical whereas the truth-values differ in different contexts. They claim, for example, that the “standards for knowledge” or “what counts as knowing” can vary with the context. That suggests that the truth-value of a knowledge attribution can vary with the context, while its content remains fixed. But such a position would be incoherent. For contextualism entails that the term “know” has variable content like other context-sensitive terms. We can understand this either as the claim that “know” expresses different relations in different contexts or that it expresses a single relation that is relativized to a contextually variable epistemic standard. Regardless whether we treat the variable content of a knowledge-ascribing sentence as indexical or as a three place relationship,⁸ a sentence of the form “*R* knows that *P*” expresses a complete proposition *only* relative to a standard, which is determined by the context.⁹ Cohen is explicit regarding this issue: “[...] strictly speaking, instead of saying that *S* knows in one context and fails to know in another, one should really say that ‘*S* knows that *P*’ is true in one context and false in the other” (Cohen, 1999, p. 65). Therefore, contextualism should be understood as claiming that its truth-value can vary only because the content of the knowledge describing sentences can vary, not that *identical* knowledge attributions made in different contexts can differ in truth-value.¹⁰ But this implies that there is *in fact no contradiction* between *C*’s and *S*’s claims. Apart from containing the same words, the proposition semantically expressed by *C*’s utterance is not the denial of the proposition semantically expressed by *S*’s utterance. Thus, the intuition of contradiction to be maintained cannot be that *S* and *C* in fact contradict one another.

Nonetheless, the most powerful consideration in favor of contextualism is that it gives us an elegant explanation of why we first *get the impression* that the sentences are contradictory. In fact, contextualism’s core advantage is to offer an explanation why skeptical arguments seem convincing but nevertheless fail. According to contextualism, we get the impression of contradiction. Yet, the commonsensian does not “lose knowledge” when the skeptic raises the

standards.¹¹ People merely fail to recognize the shift in the content of the knowledge ascriptions and thereby sense a contradiction that is not in fact there. Thus, both partners erroneously *believe* they contradict each other. Because the context sensitivity of terms like “know” is hard to see even for competent speakers, people are frequently unaware of differences in the contents of knowledge attributions. Therefore, let us try a second version of the intuition of contradiction: Perhaps the intuition is that both partners *believe* that they contradict one another? This seems true. They have a dispute over the question of whether the knowledge ascription is true or false. But if the intuition we want our solution to harmonize with is that *both believe themselves to have expressed contradictory propositions*, we run also into a problem: in a sense all presented views respect this intuition. The pure fact that *S* and *C* dispute whether *R* knows that *P* is sufficient evidence for that. On the other hand, can both participants really misleadingly sense a contradiction for the whole conversation due to the unrecognized context-sensitivity of the term “know”? After all, they explicitly discuss which conditions and standards have to be met to ascribe knowledge to *R*. Both have to be aware, at least after several standard-changing maneuvers of non-cooperation, that they use different standards or that there is at least a serious misunderstanding. Both can still assume their own standards are “the right ones”, as mirrored in the intuition of persisting standards. But they have clear indication that their conversational partner has different – in their view “mistaken” – standards. Even in case of an odd conversation, if both are minimally trained in conversational situations, their dispute is not about whose claim is right, but about which standards are more *appropriate*.¹² That implies that they are aware their standards differ. But if we have a thesis about the semantics of the term “know”, how can they still fail to realize both claims are not contradictory? First, we saw *the reason* for the intuition of contradiction (believing that there is a contradiction) is that the context sensitivity of terms like “know” is hard to see. Second, both discuss whose standards are more *appropriate*, therefore realizing that their ascriptions of knowledge are based on different standards. But then, how can they *still* sense a contradiction? There is no need for the intuition of contradiction any more. This reading of the intuition seems to result in a tension with the intuition of persisting individual standards again. Either the intuition is trivial and fulfilled, or it is stronger and violated in the GAP view as well as in both the different scoreboards view and the no scoreboards view. This implies that all three views – the different scoreboards view, the no score-

boards view, and the GAP view – are in the same position. The intuition that participants believe they are contradicting one another is granted (but only in a trivial reading), while in fact they are not.

What about the second advantage of the GAP view? Is it less skeptic friendly? If that is supposed to mean that in cases of divergent knowledge ascriptions, the skeptic's claim (b) has *not the truth-value "true"* – yes. If we demand the skeptic's claim is *false* – no. The apparent advantage of the "GAP" view results only from introducing a third truth-value. But the concurrent different scoreboards view (see Figure 1) concedes only that the skeptic's claim is true *according to his own high standard*, and because of the difference in standards, there is no shared scoreboard. In other words, there is an area within the critical domains of disagreement in which the skeptic's claim is true, but this area does not cover the whole area of disagreement. This seems a price we can pay and contextualist's have to pay – because we want to maintain the intuition of individual standards. If that is still too "skeptic friendly" to make contextualism attractive, as DeRose suggests, I suppose we have to reject contextualism.

8. SUMMARY

I have argued that there is a problem with the GAP view. First, it suffers from the same problems it is supposed to resolve – that at least one of the intuitions is violated. But for a convincing contextualist view, we need not only fix the basic parameters but answer the question how, in cases of disagreement and different individual standards, the common conversational score is determined. Second, I have argued that the basic parameters may be incorrectly formulated. In order to get the most promising contextualist answer, we *do not need* a solution that justifies both intuitions and states that the skeptic's claim is not true. For both intuitions are not on an equal footing. Rather, we may debunk one of the intuitions, i.e. the intuition of contradiction, because it is based on a serious misunderstanding of contextualism. That is not to say that there is no intuition of contradiction. Rather, it is to say that the intuition is misleading and only holds on a weak reading. As a consequence the GAP view has no serious advantage above the different scoreboards view and the no scoreboards view. In addition, I argued that the different scoreboards view and the no scoreboards view in fact coincide. In case this back door does not seem convincing, the last option available is even nastier for contextualists: If (1) there is no solution for the

question *Q* we started with that is compatible with both intuitions, and (2) we are not willing to give up either of the intuitions – it can be expected that (3) here lies a serious problem for contextualism in general. Regardless of our strategy regarding this issue, that does not relieve us from the duty of describing in detail how the personal standards are integrated into the conversational score. Closely related tasks are proposing answers to whether context changes are governed exclusively by conversational criteria, how these criteria are rated in the personal and shared score, and how the dynamic of context changes emerges. But these are different and highly demanding tasks.

ACKNOWLEDGMENTS

I am grateful to Elke Brendel and Mylan Engel as well as to an anonymous referee for this journal for discussions and very helpful comments on an earlier version of the paper.

NOTES

- ¹ To increase clarity, let us assume that there are no other vague or context-sensitive terms involved.
- ² In discussion Mylan Engel pointed out that there is another option: The reasonable standard prevails. But this view meets the same problem. The intuition of persisting individual standards is violated as well.
- ³ It is tempting to assume that (b) is in fact no possibility, because *S*'s standards are stricter than *C*'s. But because we look for a unified solution, we should not rule out the possibility that the *S*'s and *C*'s standards are not overlapping, nonetheless *S*'s standards are stricter than *C*'s. Therefore in principle a situation is possible where *S*'s standards are fulfilled but *C*'s not.
- ⁴ Lewis mentions the fact that raising the standards seems much easier than lowering them. Nonetheless they can be lowered in a conversation with high standards, according to Lewis when "something is said that is true under the lowered standards, and nobody rejects, then indeed the standards are shifted down" (Lewis, 1979, p. 245). Yet the content can seem "imperfectly acceptable". In contrast raising standards is quite easy: simply mentioning "aloud or even in silent thought" (Lewis, 1996/2000, p. 379) any particular possibility will raise the standard. Even if raising the standards interferes with our conversational purpose, it is "commendable". Regarding Lewis's version of conversational contextualism in terms of available alternatives the general question is: Which possibilities may be *properly* ignored? The hard problem is to specify the vague term *properly*. We have already mentioned the most important mechanism for conversational contextualism in relevant alternatives theories: conversational mechanisms, by which alternatives are introduced by the attributor. One simple

way of changing standards is mentioning something that would be not acceptable if standards are not changed. If nobody rejects, you bring additional alternatives into play, and by doing so, you change the standard. However, by Lewis's Rule of Attention, regardless how far-fetched a possibility is, once we are attending to it, it is a relevant alternative: "a possibility not ignored at all is ipso facto not properly ignored" (1996, p. 559). In effect, that states that no consciously recognized possibility *could* be properly ignored. As a result, a situation like the one DeRose describes cannot happen. It is just not an option in a conversation to stubbornly resist ignoring skeptic possibilities that come into play. Lowering the standards and insisting on the lowered standards when the standards were already high seems even harder – because in the process of conversation we are already attending to these possibilities.

- ⁵ Cohen dubs them as "speaker sensitive" in contrast to "subject sensitive" ones (Cohen, 2000).
- ⁶ At least regarding classical context-sensitive terms like indexicals, it is an exceptional situation, and DeRose's claim is precisely that we can treat "know" as a gradable expression as well and should understand skeptical doubts correspondingly.
- ⁷ This is far from being uncontroversial. For a convincing argument that contextualist's attempt to treat "know" as a gradable expression fail, see Stanley (forthcoming).
- ⁸ "*R* knows *C*₁ that *P*" and "*R* knows relative to *C*₁ that *P*," respectively.
- ⁹ How the context determines the standard, as Cohen observes, is a complex function of speaker intentions, listener expectations, presuppositions of the conversation, salience relations, etc. (1999, p. 61) – in other words Lewis's conversational score.
- ¹⁰ For a more detailed discussion, see Bach (forthcoming), Stanley (2004), or Williamson (2004).
- ¹¹ For an objection, see Brendel (2003).
- ¹² Otherwise not only *some* conversational rules are violated.

REFERENCES

- Bach, K.: forthcoming, 'The Emperor's New 'Knows'', in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy: On Epistemology, Language and Truth*, Oxford University Press, Oxford. Cited from: userwww.sfsu.edu/~kbach/contextualism.pdf.
- Brendel, E.: 2003, 'Was Kontextualisten nicht wissen', *Deutsche Zeitschrift für Philosophie* **51**, 1015–1032.
- Cohen, S.: 1999, 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives* **13**, 57–89.
- Cohen, S.: 2000, 'Contextualist Solutions to Epistemological Problems: Skepticism, Gettier, and the Lottery', in E. Sosa and J. Kim (eds.), *Epistemology. An Anthology*, Blackwell, Cambridge, MA, 517–530.
- Feldman, R.: 2004, 'Comments on DeRose's "Single Scoreboard Semantics"', *Philosophical Studies*, **119**, 23–33.
- DeRose, K.: 2004, 'Single Scoreboard Semantics', *Philosophical Studies*, **119**, 1–21.

- 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. Reprint 2000 in E. Sosa and J. Kim (eds.), *Epistemology. An Anthology*, Blackwell, Cambridge, MA, 503–516.
- Nichols, S., S. Stich, and J. Weinberg: 2003, 'Metaskepticism: Meditations in Ethno-Epistemology', in S. Luper (ed.), *The Sceptics. Contemporary Essays*, Ashgate Publishing, Aldershot, UK, Ashgate Epistemology and Mind Series.
- Stanley, J.: 2004, 'On the Linguistic Basis for Contextualism', *Philosophical Studies* **119**, 119–146.
- Williams, M.: 2000, 'Is Contextualism Statable?', *Philosophical Issues* **10**, Skepticism, ed. E. Sosa and E. Villanueva, 80–85.
- Williamson, T.: forthcoming, 'Knowledge and Scepticism', in F. Jackson and M. Smith (eds.), *The Oxford Handbook of Contemporary Analytic Philosophy*, Oxford University Press, Oxford.

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

MICHAEL WILLIAMS

KNOWLEDGE, REFLECTION AND SCEPTICAL HYPOTHESES

1. TWO KINDS OF CONTEXTUALISM

Keith DeRose defines epistemological contextualism as “the position that the truth-conditions of knowledge ascribing and denying sentences (sentences of the form “*S* knows that *P*” and “*S* doesn’t know that *P*” and related variants of such sentences) vary in certain ways according to the context in which they are uttered” (DeRose, 1999, p. 187). As he notes, there are several competing versions of this general view. For example subject contextualism and attributor contextualism divide over whether the shifting standards that a person must meet to count as knowing are set by the context of that person or by that of whoever is describing him as knowing or not knowing. De Rose argues forcefully in favour of the attributor (DeRose, 1999, pp. 190–191). However, subject and attributor are closely related. According to both views, the standards for (truly) saying of a person that (s)he knows that *P* can be more or less severe depending on the (subject’s or attributor’s) *conversational* context. Thus both are versions of *conversational contextualism*. Indeed, both are versions of what I call *simple conversational contextualism* (SCC). I shall say more about what I mean by this in the next section, where I articulate the position in more detail.

Conversation contextualism contrasts importantly with what DeRose calls *structural contextualism* (DeRose, 1999, p. 190). On this view, hints of which can be found in Austin and Wittgenstein, justification (hence knowledge) presupposes a definite *issue context*. Proponents of structural contextualism tend to see it as an alternative to foundationalism and coherentism. DeRose thinks that it is better seen as a (non-standard) version of foundationalism. Either way, structural contextualism must be distinguished from SCC, since the latter implies *no* view about the structure of knowledge or justification. I have my doubts about this. Accordingly, I shall adopt more neutral terminology and call DeRose’s structural contextualism *issue contextualism*.



Do these two forms of contextualism have more than a name in common? I think they do.

First, both sorts of contextualism are committed to the following generic claim about knowledge:

- (C) The standards for (truly) attributing (or claiming) knowledge are not fixed but vary (somehow) with the context in which knowledge is attributed (or claimed).

Furthermore, both exploit C in their approach to scepticism. Indeed, they offer competing articulations of the Basic Contextualist Diagnosis:

- (BCD) The sceptic reaches his paradoxical results by exploiting the context-sensitivity of epistemic standards. Sceptical conclusions seem plausible because the very practice of sceptical argumentation or “doing epistemology” tends to set epistemic standards so as to make such conclusions true. However, this fact does not invalidate everyday knowledge-claims and attributions, which remain true at everyday (“non-philosophical”) standards.

If BCD is correct, the sceptic’s mistake is to think that he has discovered, while doing epistemology, that knowledge is impossible. In fact, the *most* that he has discovered is that knowledge is impossible while doing epistemology.

This is an attractive thought. Scepticism is a problem because while, on the one hand, sceptical conclusions are difficult (if not impossible) to accept, sceptical arguments seem (or can be made to seem) intuitively compelling. BCD accounts for this smoothly. As Hume saw, while scepticism leaves us cold in everyday situations, it tends to triumph in the study (the site of “doing epistemology”). But as Hume did not see, the context-bound appeal of sceptical claims reflects the logic of knowledge-attributions, not the psychology of belief. This isn’t quite right, for this way of putting things is too concessive for issue contextualists. More of this as my argument proceeds.

So much for common ground. The topic of this paper is the differences in how attributor contextualists and issue contextualists flesh out BCD. I will explore this issue by focussing on a particular question. According to these different articulations of contextualism, why do we take sceptical arguments seriously? In particular, why do we take seriously sceptical hypotheses (that I am a victim of an Evil Deceiver, or a brain in a vat), given that all of us regard them as completely outlandish? I shall argue that issue contextualism has a much better answer to this question. As a result, issue contextualism offers an anti-sceptical strategy that promises much deeper insights into how scepticism arises and how it can be avoided.

2. SCC AND SCEPTICISM

I focus on sceptical hypotheses because SCC is typically applied to the diagnosis of Cartesian scepticism; and Cartesian scepticism is distinguished from Agrippan (regress) scepticism by the use it makes of sceptical hypotheses.

Sceptical hypotheses are a special kind of “defeater” to ordinary knowledge-claims. They are defeaters because they posit situations such that, if they obtained, our ordinary beliefs would not amount to knowledge, either because our beliefs would be false or, even if true, epistemically defective. They are special in that they involve *systematic* error or deception. Thus in both Demon deception and brain-in-vat cases, our experience is manipulated to mimic the experience we have in (what we take to be) our “normal” world. Because such scenarios incorporate systematic deception, it can seem hard to say how we could know that they do not obtain.

A standard deployment of sceptical hypotheses is found in what Keith DeRose calls the Argument from Ignorance (DeRose, 1995, p. 183). Let *O* be some ordinary claim – e.g the claim that I have two hands – and *H* be some appropriate sceptical hypothesis. The sceptic now argues:

- (AI) I don’t know that not-*H*.
 If I don’t know that not-*H*, I don’t know that *O*.
 So: I don’t know that *O*.

Since *O* could be any ordinary claim about an external object, it seems that we have no knowledge of external things.

Advocates of SCC treat AI, or some obvious variant, as the canonical form of sceptical arguments. Since SCC comes in several versions, reflecting differences over how to understand knowledge, proponents of SCC (SC contextualists) do not all say exactly the same thing about AI. Still, there are certain shared commitments, which I state in terms of SCC’s “attributor” version:

- (SCC1) “Context” is first and foremost
 conversational or *dialectical* context. That is to say, the
 standards for attributing (or claiming) knowledge
 depend on what explicit knowledge-claim has been made,
 or on what error-possibilities have been brought up,
 are being attended to, or are otherwise “salient”.

- (SCC2) The more “remote” (far-fetched, improbable) the error-possibilities (“defeaters” to our knowledge-claim) we feel obliged to consider, the harder it is to have knowledge (or more precisely, the more demanding the standards for truly attributing knowledge).
- (SCC3) Everyday epistemic contexts are *restricted*. That is to say, the error-possibilities in play are limited to a restricted range of relevant alternatives to what we claim to be the case, or presuppose in so-claiming. The restricted character of everyday epistemic contexts keeps epistemic standards (comparatively) low.
- (SCC4) The epistemic context created by doing epistemology, is *unrestricted*. When reflecting philosophically, we are open to any coherent error-possibility. Thus the effect of “going philosophical” is thus to raise the standards for attributing knowledge to the maximum.
- (SCC5) Removing all restrictions on relevant defeaters allows us to bring sceptical hypotheses into play. Alternatively, bringing such hypotheses into play has the effect of lifting all everyday restrictions. Thus, simply presenting AI creates, or tends to create, context in which knowledge-claims about ordinary objects turn false

Of course, what the proponent of this diagnosis gives with one hand he takes away with the other. Knowledge-attributions turn false in the widespread way that sceptics claim only in the peculiar context created by considering sceptical hypotheses. Ordinary knowledge-attributions made in ordinary contexts are safe from sceptical undermining.

My reasons for talking about simple conversational contextualism should be clear. Calling this form of contextualism “conversational” reflects its stress on the role of conversational context in fixing epistemic standards. Calling it “simple” underlines the idea

that there is a measure of severity for epistemic standards that applies independently of subject-matter. With respect to epistemic standards, context fixes their degree of severity on some *context-independent* scale.

This latter idea is very important. In my view, it is what distinguishes SCC from issue contextualism. Issue contextualists tend to think that epistemic standards are subject-matter sensitive. So, for example, both scientific experiments and historical researches can be conducted according to more or less strict standards. But there need be no answer to the question of whether knowledge in physics is subject to stricter standards than historical knowledge. Physics and history could be too disparate for any such comparison to make good sense. (I am talking here about physics and history as such: in either discipline, standards could deteriorate so that, at a given time, one subject was in a poorer state than the other. A subject could fail to maintain its own standards.) Noticing this feature, some philosophers associate issue contextualism with relativism. I am inclined to demur, though I will not pursue the issue. My interest here is scepticism.¹

The contrast between SCC and issue contextualism with respect to whether we should see epistemic standards as varying across contexts in ways that go beyond anything that could be captured by a single severity scale goes to the heart of the differences between their approaches to scepticism. SC contextualists think that the sceptic *raises the standards* for knowing. But issue contextualism opens another possibility: that the sceptic changes the subject; that “doing epistemology” involves a radical break with ordinary epistemic standards, so that sceptical doubting is *not* an extension of ordinary doubting but at best a partial and deceptive simulacrum of it.

Here is one reason why the difference matters. According to SC contextualists, the sceptic seems right because, in a limited way, he is right. Doing epistemology results in a simple failure of knowledge. The sceptic’s *only* mistake is to think that he has shown more than he has. The anti-sceptical strategy on offer is thus one of *pure insulation*. Such a strategy is seriously *concessive*.

A contextualist response to scepticism does not have to be concessive in this way. Identifying scepticism’s context-bound character might be only the first step, the second being to call into question the theoretical tenability of the sceptic’s context. The point of connecting scepticism with “doing epistemology” might be to raise about whether epistemology (so conceived) is worth doing. One way it might prove not to be worth doing is this: doing it may

depend on implicit theoretical presuppositions (presuppositions belonging to what I shall eventually identify as epistemology's disciplinary meta-context) that we have no reason to accept, and even many reasons no to accept. This is what issue contextualists argue, or ought to argue.

SC contextualists do not go down this road. This is not an oversight. The essence of their view is that the sceptic generates his conclusions by exploiting quite ordinary context-shifting mechanisms. To be sure, he exploits them in a surprising way, but there is nothing wrong with that. Sceptical conclusions *seem* intuitive because, within their proper limits, they *are* intuitive. Once the possibility that I am a brain in a vat is visibly on the table, I have no way of knowing that it does not obtain.

Still, a piece of the puzzle is missing. Why do we bring up sceptical hypotheses in the first place? And even if they cross our minds, why do we take them seriously? Certainly, lots of epistemologists take them very seriously indeed, at least in the sense of seeing them as having great theoretical interest. Yet none of us believes that any sceptical hypothesis is true, or even remotely likely to be true. So what is the source of their interest? No diagnosis of scepticism will be satisfactory if it leaves us in the dark on this fundamental point.

Someone might say that taking an interest in scepticism (hence in sceptical hypotheses) just is (an important aspect of) doing epistemology. In epistemology, we try to understand how knowledge is possible. But there is a question about how knowledge is possible only because there are intuitively plausible arguments for the conclusion that knowledge is *impossible*. Of course, no one is under any obligation to be interested in epistemology. If you find epistemology uninteresting, don't do it. Then you can ignore sceptical hypotheses. But if you do take an interest in epistemology, you can't.

I have some sympathy with this: certainly, concern with scepticism has been a driving force behind much epistemological theorizing. But is this really all there is to say about epistemology: that it intrigues some people and not others? I don't think so.

3. IGNORING THE PROBLEM

SC contextualists tend to take interest in sceptical arguments (involving sceptical hypotheses) as a given. Accordingly, while they have a lot to say about how sceptical arguments work, they have little

to say about why sceptical hypotheses are worth taking seriously. This is not an accident: their position lacks the necessary resources. But first let me illustrate the phenomenon.

First David Lewis. According to Lewis, knowledge is infallible in the sense that, for me to know that *P*, my evidence must eliminate every possibility in which not-*P*. (Lewis understands “evidence” and “eliminate” very broadly, so his account of knowledge is externalist in spirit.) However, depending on context, some possibilities may be ignored, so that the standards for “infallibility” are more severe in some contexts than others.

In Lewis’s terminology, the standards for knowing depend on our *presuppositions*, where to presuppose proposition *Q* is to ignore all possibilities in which not *Q*. Obviously, if we were free to presuppose, or ignore, whatever we like, knowledge would be far too easy to come by. But we are not free in this way. Rather, there are normative rules governing *proper* presupposition or *proper* ignoring. Lewis suggests a number of such rules. Particularly important is the *Rule of Attention*:

When we say that a possibility *is* properly ignored, we mean exactly that; we do not mean that it *could have been* properly ignored. Accordingly, a possibility that is not ignored at all is *ipso facto* not properly ignored. (Lewis, 1996, p. 230)

In the light of our current question – why should we take sceptical hypotheses seriously? – this seems a very odd rule. Since Lewis’s rules are normative – they govern what it is *permissible* to ignore – we might have thought that the relevant issue is precisely the one Lewis sets aside: i.e. could we have (properly) ignored something, even if we didn’t? Not only is the Rule of Attention silent here, it amounts to a license to ignore the question.

Here is another of Lewis’s rules, the *Rule of Belief*:

A possibility that the subject believes to obtain is not properly ignored, whether or not he is right to so believe. Neither is one that he ought to believe to obtain – one that evidence and arguments justify him in believing – whether or not he does so believe. (Lewis, 1996, pp. 226–227)

But again, what about possibilities that we do *not* believe to obtain, or positively believe not to obtain. Absent some reason to take them seriously, are we entitled to ignore them? Not if we notice them apparently.

There is a striking – and completely unexplained – asymmetry in Lewis’s attitudes towards ignoring and attending. I *can’t* make possibilities irrelevant merely by ignoring them; but I *can* make them relevant merely by paying them some attention. This asymmetry explains Lewis’s claim that knowledge is subject to what he calls “the

sotto voce proviso”: *S* knows that *P* iff *S*’s evidence eliminates every possibility in which not *P* – *Psst*, except for those possibilities that conflict with our proper presuppositions (Lewis, 1996, p. 225). The proviso has to be *sotto voce* because, as Lewis interprets the Rule of Attention, a possibility becomes relevant if it so much as crosses one’s mind. This is why our knowledge of the world is “elusive”: it temporarily evaporates every time we so much as think of a sceptical possibility.

This is very implausible. In our normal way of talking, ignoring and noticing are not incompatible. On the contrary, ignoring is often deliberate; and one can deliberately ignore only something (or someone) one has noticed. In explaining a problem in mechanics, a teacher might say, “In this situation, we can properly ignore resistance due to friction”. It would be silly to reply, “You just brought it up; so you aren’t ignoring it; *a fortiori*, you aren’t properly ignoring it”.

What all this suggests is that Lewis’s Rule of Attention is not first drawn from reflection on everyday conversational rules and subsequently *applied* to doing epistemology. Rather, from the outset the Rule is designed to set the bar for relevance very low, allowing Lewis to finesse the question of why sceptical hypotheses have any call on our attention.

Robert Fogelin has a more “justificationist” (and much more subtle) account of knowledge. But he agrees with Lewis that the standards we impose in attributing knowledge go up and down according to the sorts of defeaters we bring into view. In Fogelin’s useful terminology, bringing up new defeaters to a knowledge-claim tends to raise the “level of scrutiny” to which that claim is subject. When sceptical defeaters are in play, the level rises so high that we find ourselves reluctant to claim any knowledge at all.

So again, why do we bring sceptical defeaters into play. And even if we bring them into play to the extent that we notice them, what makes them relevant? To his credit, Fogelin acknowledges that there is an issue here. As he notes, Descartes, who seems to be the first to have trafficked in sceptical scenarios, had a methodological interest in them. Wanting an absolutely secure basis for knowledge, he hoped to use sceptical thought-experiments to filter out foundational certainties. But Fogelin asks: Given the failure of Descartes’s project, why do other philosophers concern themselves with sceptical scenarios, especially today? An enormous amount has been written on the topic of late: what drives this interest? Fogelin admits that he is not sure

that he knows the answer to this question. But he has a suggestion to make:

Part of the answer, I think, is this: Dwelling on remote defeaters can itself raise the level of scrutiny. Put differently, defeaters that are not salient in everyday life can be made salient simply through intensely reflecting on them. (Fogelin, 2003, p. 108)

I don't much care for this talk of salience, which Fogelin shares with Stewart Cohen, for it tends to blur the distinction between the psychological question of whether something *is* on our minds with the normative-epistemological question of whether it *ought* to be, or whether we would be offending against any important epistemic norms if, having once noticed it, we decided to ignore it anyway. True, in Fogelin's view, it is not enough merely to notice a sceptical possibility: one has to dwell or reflect intensely on it. But why does mere psychological salience establish epistemic relevance?

Fogelin's account of knowledge helps him here. For Fogelin, an important element in the justification required for knowledge is epistemic responsibility: we are disinclined to count as instances of knowledge beliefs that are held or formed irresponsibly (for example, in the teeth of counter-evidence). So if, for whatever reason, you find yourself taking a defeater to a given belief seriously, then to continue to hold that belief, while unable to cope with the defeater, is to be epistemically irresponsible. This is why we are reluctant to claim knowledge, once sceptical scenarios are (rightly or wrongly) in play.

Fogelin's view is a clear improvement on Lewis's. Even so, it is not even part of the answer to the question he poses. At best, Fogelin explains how sceptical possibilities compromise knowledge, if we take them seriously. But why do we take them seriously? There seems to be no answer to this question. Some people are gripped by scepticism, some aren't. Fogelin again:

When doing philosophy, one can be made to feel the force of Cartesian doubt. There are, it seems, certain philosophers who do not rise to the bait dangled by sceptical scenarios – either because they see danger in the offing or because they simply do not get it. They, perhaps, are blessed. (Fogelin, 2003, p. 108).

But there is an ambiguity in the phrase "Cartesian doubt", which can refer either to Descartes's standards for doubting – anything that is not absolutely certain is subject to Cartesian doubt – or to the principle target of his doubts in the *First Meditation* – our knowledge of the external world. Fogelin starts from the first use but, in the quotation just given, slides to the second. He has no explanation for why we continue to be interested in Cartesian scenarios – hence external world scepticism – given that we have abandoned

Cartesian aspirations. He says that philosophers feel the force of Cartesian doubt. But if they have given up on Cartesian certainty, the force of the doubt cannot be that nothing is absolutely certain. He says that other philosophers see danger. But what danger: that the quest for certainty is doomed? They know that already. In giving his answer, Fogelin loses track of his question.

For a final example, I turn to Keith DeRose. DeRose develops a variant of Nozick's subjunctive conditional analysis of knowledge. In determining whether *S*'s belief that *P* amounts to knowledge, we have to consider whether it matches the fact of the matter, not just in the actual world, but in all *sufficiently close* possible worlds as well. The greater the distance one can stray from the actual world, such that *S*'s belief continues to match the facts, the *stronger S's epistemic position* with respect to *P*. Knowledge is true belief involving a proposition with respect to which one stands in a sufficiently strong epistemic position.

This notion of strength of epistemic position is related to but distinct from Nozick's idea of sensitivity. A belief that *P* is sensitive given that, if it were not the case that *P*, *S* would not believe that *P*. For mundane propositions, strength and sensitivity go together: if I am in even a minimally strong epistemic position with respect to such propositions, my beliefs regarding them will also be sensitive. Last year, I went to Italy for my family vacation. Naturally, all sorts of familiar mishaps might have disrupted the trip; but if they had, I would not believe myself to have taken an Italian holiday. By contrast, my belief that I am not a brain in a vat, though it will match the facts over a wide range of situations, some quite distant from reality, is not sensitive; for I would continue to believe that I am not a brain in a vat, even if I were. So while I am in a *strong* position with respect to believing that sceptical hypotheses are false, this is only because the possibilities they raise are *so* remote that there aren't any nearby worlds in which such hypotheses are true.

Although DeRose does not build sensitivity into his analysis of knowledge, he appeals to sensitivity to explain the plausibility of the sceptic's claim that I do not know that any sceptical hypothesis is false. We have, he claims a strong, though not exceptionless, reluctance to count insensitive beliefs as knowledge. Accordingly, he proposes a Lewis-type *Rule of Sensitivity*, governing the contextually appropriate standards for knowledge:

When it is asserted that some subject *S* knows (or does not know) some proposition *P*, the standards for knowledge (the standards for how good an epistemic position one must be in to count as knowing) tend to be raised, if need be, to such a level as to

require *S*'s belief in that particular *P* to be sensitive for it to count as knowledge. (DeRose, 1995, p. 205)

Take my claim to know that I went to Italy for my vacation: the Rule will demand sensitivity for my belief and this requirement will be met across an appropriate range of worlds. Now, since I am not a brain in a vat in any world within that range, my insensitive belief that I am not will also match the facts. Accordingly, I will be in as strong a position with respect to my anti-sceptical opinion as I am with respect to what I believe about my recent travels, so that *both* beliefs amount to knowledge at the standards enforced by a mundane knowledge-claim or attribution. However, if I claim explicitly to know that I am not a brain in a vat, the Rule will require sensitivity for that particular belief, vastly expanding the range of worlds over which, to count as knowledge, *any* belief of mine must match the facts. Thus my explicit anti-sceptical knowledge-claim creates a context in which I must be in a much stronger than normal epistemic position in order truly to claim to know *anything*: even where I went for my holidays. In this way, the AI's second premise – if I know that *O*, then I know that not-*H* – is true at every level. And its first premise – that I do not know that not-*H* – is made true by an explicit knowledge-claim or attribution with respect to not-*H*. The plausibility of the argument is thus accounted for. But since everyday knowledge claims remain true by everyday standards, a blanket sceptical conclusion is resisted.

This view, too, is much more plausible than Lewis's. For DeRose, sceptical possibilities are brought into play by *explicit knowledge-claims*, not by someone's merely mentioning or noticing them. Even so, DeRose does not really explain why we take sceptical hypotheses seriously. To bring the possibility into conversational play, I need to make an explicitly anti-sceptical knowledge-claim: I have to claim to know that I am not a brain in a vat. But I will only make such a claim if I am *already* disposed to take sceptical hypotheses seriously. The source of this disposition remains to be discovered.

Like Fogelin, DeRose recognizes that not everyone resonates to sceptical arguments. Commenting on Barry Stroud's claim that Descartes's "dreaming" argument "appeals to something deep in our nature and seems to raise a real problem about the human condition", DeRose notes that some people have a quite different reaction, encapsulated by the exclamation "Aw, come on".² He takes this reaction to imply a judgement that the sceptical argument is extremely weak. His reply is that it isn't extremely weak but rather strong: it is

clearly valid and, considered individually, its premises “enjoy a good deal of intuitive support”. You may think you know that you are not a brain in a vat, but how *could* you know such a thing?³

This response anticipates the concessive character of the response to scepticism DeRose favours, but that is about all. Notice that DeRose’s rhetorical question implicitly excludes the Moorean response that ordinary knowledge excludes sceptical possibilities: knowing that I have hands, I know that I am not a brain in a vat. It also excludes the Quinean variant of Moore’s response: that science excludes them. Maybe it’s *logically* possible that I’m a brain in a vat, but it’s not *technologically* possible. In posing his apparently innocent question – How *could* you know that you are not a brain in a vat? – DeRose has already projected himself into a context in which ordinary (and scientific) knowledge has been put up for grabs. How does this happen? The answer is: by virtue of our *already* being determined to take sceptical scenarios seriously. Why?

DeRose’s answer seems to be Fogelin’s: some of us are just built that way. However, for DeRose, who is very sensitive to questions of burden of proof, this is a problematic move. As he notes, his own response to scepticism involves assuming things that the sceptic claims we can’t know. For example, in claiming that his belief that he has hands is sensitive, he betrays his conviction that he is not a brain in a vat, in this or in any nearby worlds. Is it legitimate to use this conviction against the sceptic? According to DeRose, the answer is “No”, if we are playing King of the Mountain: that is, if we are trying to prove, in terms acceptable to the sceptic, that scepticism is false. But this cuts both ways, so that

...if the sceptic is marshalling deeply-felt intuitions of ours in an attempt to give us good reasons for accepting his skepticism, it’s legitimate to point out that other of our beliefs militate against his position, and ask why we should give credence to just those that favour him. (DeRose, 1995, p. 215)

While this is a fair question, it can just as well be posed by the “Aw, come on” school of Moore and Quine. If our aim is not to refute the sceptic, once he has been handed all the cards, why take the long way round? Why not use our commonsense convictions to head him off right at the beginning? Particularly, why not do this, given that we have no interest in Cartesian certainty?

One reply is that the aim is not to *refute* the sceptic but to *understand* him. A good response to scepticism should be diagnostic and not merely dialectical. After all, the sceptic is not so much an opponent as a personification of our own tendencies to be swayed by

sceptical arguments. Moore's way with the sceptic is thus not so much wrong as unenlightening. It gives us no insight into why we find scepticism so plausible.

While I am sympathetic to this reply, it takes a lot for granted. DeRose's talk of "deeply felt intuitions" is reminiscent of Stroud's claim that scepticism appeals to something "deep in our nature". DeRose never so much as considers the possibility that our intuitions – which are the intuitions of professional philosophers – are just an artifact of our philosophical education. Still less does he consider that the possibility this education itself involves initiation into a tradition that has long since slipped into a degenerate, scholastic phase. DeRose wants a partially vindicating explanation of scepticism's intuitive appeal, one that reveals the sceptic as subtly misusing intuitions about knowledge that deserve to be endorsed, rather than explained away. There is nothing wrong with looking for such an explanation. But we should remember that such a strategy encourages us to think that the sceptic must be right about something. In consequence, it has an inbuilt tendency to favour responses to scepticism that are to some degree concessive. As we noted SCC is seriously concessive. We should not be surprised.

I will not press this line further. Not that I think that there is nothing to it: quite the contrary. However, there are reasons to suspect that it can't be the whole story. Most philosophers think that philosophical scepticism turns radical. Sceptical arguments seem to show not just that we fail to have knowledge by exalted Cartesian standards, but that we fail even by relaxed, everyday standards. This is the answer to Fogelin's question: "Why are we still interested in Cartesian (external world) scepticism?" How does this radical turn even *seem* to happen?

The issue of diagnostic adequacy is crucial. But that has always been my point. The fact that SC contextualists have so little to say about why sceptical possibilities are interesting (beyond the fact that they are useful in arguing for scepticism) is significant because it points to the inadequacy of their diagnosis.

4. REFLECTION ALONE

SC contextualists think that the sceptic makes an unusual (and perhaps deceptive) use of ordinary context-shifting mechanisms. But as Fogelin insists, sceptical arguments presume that the level of scrutiny can be raised by reflection alone. Does this ever happen in ordinary

cases? Or is the sceptic's attempt to change standards peculiar in and of itself?

Let us consider an ordinary example of raising standards by introducing a new defeater.

Timetable. I have an appointment that I cannot afford to miss; but I also have important things to do beforehand. I therefore ask you if you know when the last train leaves that will get me to the city in time. You say you do. Needing to be reassured, I ask you how you know. You say you happen to have just consulted the timetable and offer to show it to me. I notice that it is last year's edition. Have you looked into whether there have been any revisions? No. So your timetable could be out of date. Perhaps you don't know when the train leaves. Certainly, we have some inclination to judge that you don't.

In this case, which I think is quite typical, we find three factors that contribute to raising the level of scrutiny. Two are epistemic (having to do with reasons and evidence) and one is "economic" (having to do with the costs and benefits of getting things right or wrong).

Red Flag. A specific piece of information suggests a particular error-possibility. Here, the fact that the timetable is not current, suggests that it may not be accurate.

Background Information. An ostensible red flag could be neutralized by knowledge that rules out the error-possibility it points to. Relative to our background information, there must be *some* likelihood of that possibility's being realized: the indicated error-possibility must be *live*. Here, I know that timetables do get revised, and I don't know that it hasn't happened in the case at hand.

High (enough) stakes. Even flagging a live error-possibility may not be enough to make it worth taking seriously. The possibility may be extremely remote and the costs of error low. But if a mistake has serious consequences, even a rather unlikely error-possibility, once flagged, can be worth considering. This is what is going here. While I accept that the chances of the timetable's having been changed are not high, I simply can't afford to miss my train.

In sum, we raise the level of scrutiny by flagging an error-possibility that, relative to our background information and stakes (accounting for opportunity and information-gathering costs), is probable enough to be worth considering.

Red flags are not always necessary, as we can see by modifying the timetable case:

Track repairs. Your timetable is up-to-date. But it occurs to me that, from time to time, trains are delayed because of track repairs. Do you know that no such repairs are taking place today? No. So perhaps you don't know when the last suitable train leaves.

Here there is no red flag. But given appropriate stakes, our background knowledge alone can certify an error-possibility as sufficiently probable. This will happen readily when the probability is reasonably high. However, the fact that error possibilities can become salient without red flags does not show that the level of scrutiny can be raised by reflection alone.

Can attention alone make error-possibilities relevant, at least in any ordinary case? Consider Gilbert Harman's lottery puzzle:

Lottery. I buy one ticket out of n (a very large number) sold. The probability that I lose is $1 - 1/n$, which in a large lottery is very close to 1. But even though I recognize that the chance of my winning is very remote, it does not seem correct to say that I know that I will lose, if this statistical information is all I have to go on. On the other hand, we do not hesitate to see ourselves as gaining knowledge from sources that are less than 100% reliable: for example, the testimony of a trustworthy informant. Yet the probability that what I come to believe is true, given testimonial evidence, may be lower than the probability that I lose in the lottery.

Stewart Cohen has suggested that we can explain this apparent discrepancy in terms of the ways in which the character of our reasons can call our attention to error-possibilities that, in other circumstances, would remain unnoticed. The explanation for our reluctance to grant knowledge in the lottery case

...lies in the statistical nature of the reasons. Although, as fallibilists, we allow that S can know q , even though there is a chance of error... , when the chance of error is salient, we are reluctant to attribute knowledge. Statistical reasons of the sort that S possesses in the lottery case make the chance of error salient. The specification that S 's reason is the $n-1/n$ probability that the ticket loses, calls attention to the $1/n$ probability that the ticket wins. Our attention is focused on the alternative that the ticket wins and this creates a context in which we are reluctant to attribute knowledge, unless S has some independent ground sufficient for denying the alternative. (Cohen, 1988, p. 106)

Is this a case in which mere attention creates salience, and hence relevance? I don't think so. The lottery case is the track-repair case all over

again: i.e. a case in which an error-possibility becomes salient *via* the interaction of background information and practical interests? When we enter a lottery, in the hope of winning a large sum of money, we decide then and there that nothing will count as conclusive evidence of losing short of the result's being officially announced. Otherwise, if we know that we are going to lose, why buy the ticket, or why not throw it away? Reflecting on the character of our reasons for thinking we will lose doesn't raise the level of scrutiny. It has been raised already.

Some evidence for this suggestion can be found (at least according to my intuitions) in the asymmetry between the first and third person cases. Imagining myself as a lottery participant, I find myself reluctant to judge that I know that I will lose. But imagining myself reacting to my friend's constantly playing the state lottery, I think I know that he is wasting his money. True, Cohen presents a third-person case. But he does so in a way that invites us to focus on *S*'s reasons. In effect, he invites us to put ourselves in *S*'s shoes. Putting ourselves in *S*'s practical situation is what elicits the reaction that we don't know that *S* will lose.

If this is right, we do not need to get involved with scepticism in order to understand how the level of scrutiny gets raised in ordinary situations. If anything, examination of ordinary context-shifting intensifies our sense of the extraordinary character of the context supposedly created by "doing epistemology". Given the appropriate combination of background information and stakes, we can see how merely thinking of a defeater can raise epistemic standards. But this is a far cry from showing how standards can intelligibly be raised by reflection *alone*.

In sceptical reflections, standards cannot be raised in the ordinary ways just scouted. Since sceptical possibilities are designed to be (supposedly) ineliminable, we cannot have evidence for or against them, and they cannot be flagged. Since, if taken seriously, they eliminate background information, along with specific claims, we have no way of estimating the likelihood of their being realized. True, most of us believe that their probability is vanishingly small. But this belief reflects our common-sense and scientific picture of the world, which the sceptic means to put up for grabs. Finally, the notion of stakes has no clear application. If I am a brain-in-a-vat, I won't catch my train, and I won't make my appointment. But in the image I will "catch my train" and "make my appointment". So it is all the same to me, as far as I will ever know.

To sum up: attributor contextualists want sceptical doubt to involve a natural extension of ordinary doubting: they want it to be no

more than an extreme instance of raising the level of scrutiny by introducing new defeaters. But they also want sceptical defeaters to be ineliminable, so that knowledge-claims really do go false in “sceptical” contexts. This means that sceptical defeaters must be capable of being made relevant in a way that does not seem to be ordinary at all, for in ordinary cases, mere logical possibilities may properly be dismissed. If anything, reflection on ordinary situations suggests that conversational developments induce standard-shifts only in definite issue-contexts (set by stakes and background knowledge). Attributor contextualism starts to look superficial. If it has anything going for it, it is that its friendliness to sceptical hypotheses offers an illuminating diagnosis of scepticism. But as we have seen, this friendliness is not well-motivated. We need to do better and I think we can.

5. KNOWLEDGE AS SUCH

None of the foregoing entails that there could not be a reason for taking sceptical hypotheses seriously. But it would have to be a reason of a different type from any we have so far isolated.

The natural question to ask at this point is: what about purely theoretical inquiry? Here there are no external costs. The cost is simply getting things wrong. This isn’t really true: there are always opportunity costs, if only those involved in not following up alternative lines of inquiry. But I will not press this point. Questions about constraints on theoretical inquiry open up a new and fruitful line of investigation.

We can agree that doing epistemology is a form of purely theoretical inquiry. It is purely theoretical inquiry into the nature of knowledge. Not any particular kind of knowledge, but knowledge as such. Does this explain the relevance of sceptical hypotheses? Fogelin thinks that perhaps it does. He writes:

why should the activity of philosophizing lead us to take cartesian skepticism seriously? Part of the reason might be that in philosophizing we are not interested in knowledge of any particular kind. We are interested in knowledge *qua* knowledge. Because of this, nothing puts constraints on the range of relevant or salient defeaters. The act of philosophizing done in a certain way makes every possible defeater salient, and, with that, skepticism is inevitable. (Fogelin, 2003, pp. 108–109)

I am not sure that I can imagine a context in which *every* possible defeater to some claim is *salient*. If everything stands out, nothing does. I take it that what Fogelin really means to say is that, in philosophizing about knowledge, there are no constraints on the range

of relevant defeaters, so that any defeater can properly be made salient by our attending to it. But this doesn't seem right either, for it leaves out something that Fogelin is well aware of, namely, the *special* interest that seems to attach to sceptical defeaters. When he first posed the question of why we take scepticism seriously, Fogelin found himself wondering about why philosophers spend so much time on sceptical scenarios. The explanation cannot be that, in philosophizing, all defeaters are equally salient. Or to put the question the other way around: why is it that, in philosophizing about knowledge, ordinary defeaters seem to be irrelevant?

Fogelin's claim that, when we take an interest in knowledge qua knowledge, there are no constraints on the range of relevant defeaters doesn't seem to be right. But I would go farther: I don't think that it *can* be right.

Some well-known remarks of Wittgenstein are very helpful.

163. ...We check the story of Napoleon, but not whether all the reports about him are based on sense-deception, forgery and the like. For whenever we test anything, we are already presupposing something that is not tested...

337. One cannot make experiments if there are not some things that one does not doubt. But that does not mean that one takes them on trust...

If I make an experiment I do not doubt the existence of the apparatus. I have plenty of doubts, but not *that*.

341. [T]he questions that we raise and our doubts depend on the fact that some propositions are exempt from doubt, are as it were like hinges on which those turn.

342. That is to say, it belongs to the logic of our scientific investigations that certain things are in deed not doubted.

343. But it isn't that the situation is like this: We just can't investigate everything, and for that reason we are forced to rest content with assumption. If I want the door to turn, the hinges must stay put.⁴

That in a given inquiry some doubts are *hors de combat* has nothing to do with either credulity or limited resources. Rather, it is a matter of the *focus* or *direction* of inquiry. What we are looking into is a function of what we are leaving alone.

In a particular discipline, there will be certain quite general pre-suppositions that serve to give that discipline its characteristic shape and subject-matter. I like to call them "methodological necessities". Together, they determine the *disciplinary meta-context* for all inquiries of a certain genre. However, they generally are not – and probably could not be – exhaustively catalogued. That is why Wittgenstein insists that certain things are *in deed* not doubted.

The point that what is and is not up for grabs determines the focus of interest applies at all levels. So particular contexts of inquiry, within a given genre, will have their peculiar local presuppositions, as will quite ordinary contexts of epistemic evaluation. But as we saw in the previous section, ordinary contexts typically involve practical concerns.

Fogelin's talk of "no constraints", like Bernard Williams's talk of philosophy as "pure inquiry",⁵ suggests that philosophizing is not only free of practical considerations but entirely presuppositionless. However, a form of inquiry that was presuppositionless would be no form of inquiry at all. Fogelin's suggestion that, in studying knowledge qua knowledge, all possible defeaters are salient is not a slip. It points to a fundamental misconception.

The question to ask is this: What presuppositions determine the disciplinary meta-context for investigating knowledge as such. Or rather, since the issue is Cartesian scepticism, what is the disciplinary meta-context for investigating the possibility of knowledge of the external world. We can approach this question by considering a specific instance of AI:

(SAI) I don't know that I am not a brain in a vat.
 If I don't know that I am not a brain in a vat,
 I don't know that I have hands.
 So: I don't know that that I have hands.

In DeRose's schematic presentation of AI, the claim that I have hands is represented by "*O*". The choice of letter is meant to suggest how the sceptical hypothesis *H* can be used to show that knowledge fails even in the case of the most "ordinary" claim. But the claim that I have hands is very far from ordinary. The oddity of the sceptical defeater is matched by the oddity of the claim it defeats. Even though I have hands and know that I do, this is not something that I would ordinarily have occasion either to claim or to claim to know. Perhaps if I were involved in some grisly accident, I might be relieved to discover and pleased to announce that I have hands. I might even assure you that I know that I have hands: I checked. But outside such outre circumstances, either claim would be distinctly odd.

There is one exception: the claim may not seem so if *we are discussing scepticism*. The claim that I have hands, or that I know that I have hands, is not so much "ordinary" as Moorean. And a Moorean claim is a set-up for the sceptic. Two essential features let it do its work.

One feature is semantic, having to do with content: i.e. with “claim” as what is claimed, what proposition asserted. A Moorean claim – and we could just as well say Cartesian claim – asserts the existence of what Stanley Cavell calls a “generic object”: the sort of thing that anyone can recognize, though without being able to say how.⁶ Contrast Moorean claims, involving generic objects, with “Austinian” claims, involving things with definite identifying features.⁷ To borrow Austin’s own example: if I say, looking at a bird, “That’s a goldfinch”, and you ask me how I know, I can reply “By its red head”. This is not the sort of thing that just anyone would know. Being able to spot goldfinches is a matter of (mildly) specialized knowledge. It follows that an Austinian claim might exemplify a particular kind of knowledge (ornithological knowledge, say), but cannot stand for knowledge as such. But this is just what Moorean claims are supposed to do. As claims involving generic objects, they are intended as generic – thus representative – claims. Reference to generic objects is a generalizing device.

The second feature is pragmatic, having to do with “claim” as an act of claiming. As we saw, considered as a speech-act, what the sceptical argument represents as an ordinary claim is really quite extraordinary. Equally clearly, the pragmatic oddity of a Moorean claiming flows from what is claimed. Typically, to assert the existence of a generic object is to give voice to something that anyone can be expected to know; and why would anyone do that? For this reason, Cavell, says that a Moorean claim is defective, by virtue of being entered in a “non-claim” context. However, that is not how the sceptic (or traditional epistemologist) sees things. To be sure, Moorean claims, entered out of the blue, are (seemingly) not tied to any special occasion of utterance, where for one reason or another there might be a question of checking up. But the sceptic sees this as effecting an essential decontextualisation. Detaching a claim from all specific contexts of utterance ensures that its epistemic appropriateness will depend *entirely* on generic epistemic factors. A Moorean claim is intended as a generic claiming. Again, the generic objects play an essential generalizing role, directing our attention to knowledge as such.

From the very outset, then, in the context of philosophical reflection on our knowledge of the world, we are trying to understand knowledge of the world in general. In a sense, the possibility of such knowledge is up for grabs before the sceptical conclusion is reached. This explains why SC contextualists join the sceptic in not backing an

“ordinary” claim against a sceptical hypothesis, as we saw in the case of DeRose.

Now let us turn to how different kinds of claims may be challenged and defeated. Austinian claims invite what we may call *criterial challenges*. The bird is a goldfinch, I claim because it has a red head. You reply: “For all that shows, it could be a goldcrest, for they have red heads too”. The challenge invokes a defeater in the form of an alternative factual possibility uneliminated by what I have offered as conclusive evidence (an identifying feature). However, this kind of challenge is as useless to the sceptic as the claim it challenges for it is as specialized as the claim it threatens to defeat. While it might show that I don’t know that the bird on the fence is a goldfinch, it has no tendency to show that I don’t know that there is a bird on the fence.

In the goldfinch case, I respond to your “How do you know” by citing an *identifying feature*. Clearly, in the case of a generic object, such a response is out of the question. Generic objects thus demand *epistemic challenges*. The suggestion must be that, in some way, I am not well-placed to make my claim: that I am under some kind of *epistemic disability*. For example;

There’s a bird on the fence.

How do you know? I can’t see anything. (So you can’t either.)

But maybe while *you* can’t see the bird (from where you are standing, it is hidden by a bush), I can. Local epistemic disabilities are of no use to the sceptic because they are in principle remediable (by an improvement in local circumstances). What the sceptic needs is an epistemic disability that is not tied to specific situations. Any disability that pervasive will of course threaten, to be irremediable. Sceptical hypotheses fill the bill. This is what makes them relevant: they belong essentially to the disciplinary meta-context of the study of knowledge as such. If introducing them raises the standards for knowing, this is an effect of their generic character. Remoteness is neither here nor there. Indeed, as we shall see in the next section, in the context of philosophical reflection on knowledge as such, it is not easy to maintain that sceptical possibilities are all that remote.

To sum up, by prescinding from specific identifying features, Moorean claims invite epistemic defeaters. But by also prescinding from particular circumstances of claiming, they invite the introduction of defeaters citing generic epistemic disabilities. Moorean claims and sceptical hypotheses are made for each other. By involving them, AI offers a paradigm for reflection on knowledge as such. But not so

fast. While I have been speaking of sceptical hypotheses as defeaters for Moorean claims, Moore thought that such claims defeat scepticism. We seem to have a standoff. And if we share Moore's intuition, why not resolve it in our favour by appealing to DeRose's point that we are not playing King of the Mountain? The answer to this is that AI does not reveal the full structure of the sceptic's thought. Sceptical hypotheses are only *indirect* defeaters to Moorean claims. They are direct challenges to a Moorean claim's *implied epistemic commitments*. This is the only way a sceptical hypothesis can suggest an epistemic *disability*.

Suppose, following Moore, I say "Here is a hand" and you ask "How do you know?". I can hardly reply, citing an identifying feature, "By the fingers". This is not because fingers are not identifying of hands: they are. The problem is that fingers are no easier or harder to identify than hands. Here we see another important function of reference to a generic object: to shift the focus of attention from the *character* of an object (which we agree is there) to its *existence*. This shift goes along with the shift from a criterial to an epistemic challenge. So if I agree to entertain your question, I have to give an epistemic response; "There's a hand here because I can see that there is".⁸

Even this will not automatically take the sceptic where he wants to go. For ordinarily taken, this "reply" may be a *dismissal* of the question. Its force may be: "I can see it (you idiot)". Or as DeRose likes to say, "Aw, come on". However, the sceptic has another way of taking this response. "I can see it" becomes "By means of the senses". Thus the sceptic takes "I can see it" to make reference to a *generic source of knowledge*, which must be presumed to be reliable, if the reply that invokes it is to secure epistemic entitlement. Naturally, he does not see himself as putting forward a tendentious re-interpretation of an ordinary response but as making explicit an epistemic presupposition involved in any claim to knowledge of the world: that the senses are a reliable source of worldly information. But am I epistemically entitled to this presupposition, given that I could be a brain in a vat?

We might think that the idea of the senses as a generic source of knowledge is innocent enough. Indeed, who could deny that there are such sources: perception, memory, testimony, and so on. However, while talk of generic sources might be innocent in itself, the sceptic (or traditional epistemologist) interprets it in a tendentious way.

Exploiting the idea of the senses as a "faculty", he encourages us to think of the senses as an information-gathering *module*. This

conception encourages two further thoughts. First, that there is a particular kind of information that the module is designed or adapted gather, so that the sensory module is *informationally bounded*. Second, that the module operates independently of other modules, so that the knowledge it produces is epistemically and semantically independent of all collateral commitments. Accordingly, we can be in possession of the evidence of the senses without being in possession of any knowledge lying outside their informational bounds. The evidence of the senses constitutes an *autonomous stratum* of knowledge. (We are by now well on the way to foundationalism, in its classical form.)

In the context of this conception of the senses, sceptical hypotheses really come into their own. If I am a brain in a vat, my senses are wildly unreliable as a source of information about the world. But surely, they would still tell me something. The thought is by now almost irresistible that they tell me how things look or appear. The evidence of the senses is *phenomenal* evidence. But since these senses are autonomous, phenomenal knowledge must be independent of knowledge of the world. Unfortunately, as sceptical hypotheses also show, phenomenal evidence radically underdetermines worldly facts. It must, since it is knowledge that we would possess even if an appropriate sceptical hypothesis were true. In the very way that they serve to set the sensory module's informational bounds, sceptical hypotheses suggest that those bounds cannot be crossed.

The disciplinary meta-context for investigating knowledge (of the world) as such is rich in presuppositions. It is therefore by no means evident that the sceptic is studying ordinary knowledge. I do not believe that he is.⁹ Rather, his very way of framing his questions involves the creation of a special subject-matter: knowledge as such. And to suppose that knowledge of the world, as such, is even a potential object of theory or reflection, we have to conceive of our epistemic capacities in a special way. The sceptic changes the subject in more ways than one.

6. A DILEMMA FOR SCC

From the standpoint of issue contextualism, SCC's chief failing is that it is not contextualist enough. Behind its contextualism is a form of invariantism: the idea that there is a simple scale by which epistemic standards can be judged relaxed or demanding, no matter what the subject at issue. But this view is difficult to maintain, even for advocates of SCC themselves.

SC contextualists want to insulate ordinary knowledge-claims from the results of sceptical reflections. But in practice, this turns out to be easier said than done.

Let us consider two further rules of conversational presupposition suggested by Lewis:

Rule of Actuality. The possibility that actually obtains is never properly ignored.

Rule of Resemblance. “[If] one possibility saliently resembles another [then] if one of them may not properly be ignored, neither may the other.”

As Lewis notes, actuality is not eliminated by the subject’s evidence. Unfortunately, any possibility uneliminated by the subject’s evidence – including sceptical possibilities – resembles actuality in this salient respect. So sceptical possibilities are never properly ignored. Result: scepticism. Lewis suggests making an *ad hoc* exception for resemblances of this type. It would, he says, be better to avoid *ad hoc*ery. But with admirable candour, he admits that he does not know how to do this.

It should be clear what to say about this. The supposed salient resemblance between our “normal” world and sceptical alternatives is that, in all worlds, the same possibilities remain uneliminated by our “evidence”. But this supposed resemblance is an artifact of the special meta-context created by the presuppositions of “traditional” (sceptical) epistemology. Outside of that meta-context, this alleged resemblance has no salience (no relevance) whatsoever. Of course, in saying this, we go beyond SCC to issue contextualism, not modifying Lewis’s approach but abandoning it.

To see that this is not just a problem for Lewis, let us turn again to DeRose. An apparent advantage of DeRose’s position, as compared with Lewis’s, is that it is more purely externalist. DeRose explains the plausibility of the sceptic’s claim that we not know not-*H* by appeal to his modified subjunctive conditionals analysis. Since his approach makes no mention of evidence’s failing to eliminate possibilities, he may seem set fair to avoid Lewis’s problem. But on closer examination, this apparent advantage proves illusory.

A well-known problem for the subjunctive conditionals analysis is given by Nozick himself in his grandmother case.

Grandmother. A grandmother learns that her grandson is alive and well when he visits her. But if he were dead or gravely ill, the family would find a way to shield her from this upsetting news. Here, SCA is violated: if the grandson were not alive and well, the

grandmother would still believe that he was. But we are reluctant to deny that she knows her grandson is alive, when she can see that he is.

DeRose agrees with Nozick: examples like this show the need to link the analysis with *methods* of belief-formation. But where Nozick favours a more complex statement of the sensitivity requirement, involving an explicit reference to methods, DeRose suggests that, in determining the range of worlds across which *S*'s belief needs to be sensitive, we should place "heavy emphasis...upon similarity with respect to the method of belief-formation utilised by *S*" (DeRose 1995, p. 196).

This move, plausible enough in its own way, threatens to saddle DeRose with Lewis's problem. In his initial presentation of his response to scepticism, DeRose's language strongly suggests that his idea of "distant" possible worlds involves a *content-based* measure: sceptical possibilities are remote in that they invoke worlds in which things happen that are wildly at variance with our ordinary view, indeed in which most of our ordinary beliefs are false. But if, in judging which worlds are relevantly close, we are to weigh similarity of methods "very heavily", it is up for grabs whether vat-worlds are remote. The sceptic claims they are not, since they resemble the actual world in respect of the role of experience in belief-formation. This is just what Lewis's Rules of Actuality and Resemblance also suggest.

DeRose is under pressure to go along with Lewis here. For not merely does he concede that the subjunctive conditionals analysis of knowledge needs to be linked with methods of belief formation, he sees, that to account for the intuitive appeal of scepticism, methods need to be individuated along the lines suggested by the sceptic.

A belief that a sceptical hypothesis is false will be insensitive. However, mere insensitivity doesn't seem to capture the appeal of scepticism. DeRose himself brings this out in an ingenious way by noticing the problem created by "naked" sceptical hypotheses: e.g. the hypothesis that I falsely believe that I have hands. While insensitive, such a hypothesis would be a very poor candidate for "*H*" in an instance of AI. Clearly, what is wrong with naked sceptical hypotheses is that they stipulate that I am wrong about some ordinary belief without saying how I came to be so mistaken. But if we ask ourselves what makes the brain-in-a-vat hypothesis so disquieting – compared with an ineffective, naked sceptical hypothesis – there is

only one answer: the envatted brain is provided with the *same experience* that an ordinary, embodied person might enjoy. But once we build this reference to common experience into our idea of method of belief formation, and once we commit ourselves to weighing similarity of methods very heavily in deciding which worlds count as “close”, sceptical worlds are brought close to the actual world. Sceptical possibilities become relevant alternatives in the most mundane situations. Result: scepticism. At the end of the day, DeRose and Lewis are in the same boat.

Again, the source of the problem is the unidimensional character of SCC, with its emphasis on raising the standards for knowledge-ascriptions. As we noticed at the very outset, this amounts to a significant invariantist element in SCC. In DeRose’s version, context determines how far out from the actual world we need to go in determining whether we are in an epistemically strong position with respect to believing that *P*. But the distance measure remains fixed. This is a bad idea. For as we just saw, in order to explain why scepticism is *ever* appealing, we have to allow that the alleged resemblances between our situation and that of the brain in a vat are sometimes salient; and then the problem is to explain in some *non-ad hoc* way why they aren’t always so. Issue contextualism provides the answer, for it allows us to argue that a shift of disciplinary meta-context doesn’t simply raise epistemic standards (on a fixed scale): it changes the subject. In DeRose’s terms, which possible worlds count as close and which as distant *also* changes with context. This is a point that DeRose himself is under pressure to concede, though to concede it is to give up on SCC altogether.

7. CONCLUSION

Issue contextualism links the apparent plausibility with the presuppositions of epistemology, conceived as the study of knowledge as such. But merely to call attention to these presuppositions is not to show that epistemology, so conceived, is in any way objectionable. Even so, the very idea of such a subject is more peculiar than is generally recognized.

Why are we inclined to suppose that it is possible to theorize about knowledge (or knowledge of the world) as such? Why do we think that we ought to be able to “understand human knowledge in general”? Barry Stroud insists that there is nothing obviously absurd about the quest for such understanding. According to Stroud, “we

can and do reflect in very general terms on human beings and their place in the world". When we do, we find both similarities and differences between human beings and "other parts of nature". On the difference side of the ledger, humans, unlike rocks and tree branches,

do not just move, they do things. Other animals also do things, but humans differ from them in the extent to which they think about what they do, and then act as a result of that deliberation. And, not co-incidentally, also know things about the world around them. Knowledge is essential to deliberation and to informed action (Stroud, 1996, p. 122).¹⁰

Such reflections lead naturally to questions about knowledge itself:

What exactly is knowledge, and how do human beings know the sorts of things they have to know to live the kind of lives they lead. These are very general questions, but they are not for that reason alone illegitimate... (Stroud, 1996, p. 123)

There is something odd about these remarks. If we want to know how human beings come to know what they need to know, to live the lives they lead today, we might have thought that, say, the history of science and technology would be the place to go for enlightenment. How could the study of "knowledge as such", which abstracts from everything having to do with human life and human interests have anything to tell us in response to the sorts of questions that Stroud raises? In truth, the crucial move is made before Stroud gets to knowledge. It is made when Stroud refers to "human beings and other parts of nature". For issue contextualists, practices of argument and inquiry do not belong to *nature* but to *culture*. Accordingly, they have particular histories rather than a common essence. Stroud's apparently commonsense reflections depend on a kind of scientism that he does not examine, or even really acknowledge.

I cannot now further contest the presuppositions of sceptical epistemology's disciplinary meta-context, though it will be obvious to anyone acquainted with the philosophy of the last century that they are eminently contestable. My point is a narrower one. The interest of sceptical hypotheses is tied to the curious meta-context we have been excavating. Thus, to the extent that we find reason to reject any or all of its presuppositions, sceptical hypotheses lose their significance. Merely mentioning them gives them no call on our attention and never did. By supposing the contrary, SCC offers only a shallow and misleading diagnosis of scepticism. If anything, it obscures the features of sceptical argumentation that we need to focus on. By contrast, issue contextualism helps us ask exactly the questions that we need to ask.

NOTES

- ¹ In Sosa/Kim (eds.) 2000, selections from my own work on issue contextualism are set in a section on relativism. I offer some thoughts about the relation of issue contextualism to relativism in Williams 2001, ch. 19.
- ² Stroud 1984, pp.39. Quoted by DeRose in DeRose and Warfield (eds.) 1999, p. 3.
- ³ Ibid.
- ⁴ Wittgenstein 1969, paragraph numbers in original.
- ⁵ Williams 1978. See especially ch. 1. For some extended criticism of Williams, see my *Unnatural Doubts*, pp. 211f.
- ⁶ Cavell 1979. See especially ch. 6. My discussion here is deeply indebted to Cavell. But I have some criticisms too. See *Unnatural Doubts*, ch. 4.
- ⁷ Austin's views, as discussed here, are found in his paper "Other Minds" in Austin 1961. For some excellent recent discussion of Austin, see Kaplan 2000.
- ⁸ A point well made by Cavell. See Cavell 1979, p. 161.
- ⁹ Neither does Kaplan. But Kaplan takes me to task for not sufficiently appreciating Austin's response to the sceptic. However, my criticism of Austin is not that he is wrong, but that he is so unsympathetic to traditional epistemology that he is not interested in excavating its disciplinary meta-context, and so not interested in understanding why doing epistemology that way might even seem to be an appealing project.
- ¹⁰ Stroud's essay is a response to my *Unnatural Doubts*.

REFERENCES

- Austin, J. L.: 1961, *Philosophical Papers*, Oxford University Press, Oxford.
- Cavell, S.: 1979, *The Claim of Reason*, Oxford University Press, Oxford.
- Cohen, S.: 1988, 'How to be a Fallibilist', *Philosophical Perspectives* 2, 91–123.
- DeRose, K.: 1995, 'Solving the Skeptical Problem', *Philosophical Review* 104, 1–52. Reprinted in K. DeRose, and T. Warfield: 1999, *Skepticism: A Contemporary Reader*, Oxford University Press, Oxford, 183–219. References to this reprinting.
- DeRose, K.: 1999, 'Contextualism: An Explanation and Defence', in J. Greco, and E. Sosa (eds.), *The Blackwell Guide to Epistemology*, Blackwell, Oxford, 187–205.
- Fogelin, R.: 2003, *Walking the Tightrope of Reason*, Oxford University Press, Oxford.
- Kaplan, M.: 2000, 'To What Must an Epistemology Be True?', *Philosophy and Phenomenological Research* 61, 279–304.
- Lewis, D.: 1996, 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549–567. Reprinted in K. DeRose and T. Warfield: 1999, *Skepticism: A Contemporary Reader*, Oxford University Press, Oxford, 220–239. References to this reprinting.
- Sosa, E. and J. Kim (eds.): 2000, *Epistemology: An Anthology*, Blackwell, Oxford.
- Stroud, B.: 1984, *The Significance of Philosophical Scepticism*, Oxford University Press, Oxford.
- Stroud, B.: 1996, 'Epistemological Reflection on Knowledge of the External World', *Philosophy and Phenomenological Research* 56, 345–358. Reprinted in B. Stroud: 2000, *Understanding Human Knowledge*, Oxford University Press, Oxford, 122–138. References to this reprinting.

- Williams, B.: 1978, *Descartes: The Project of Pure Inquiry*, Pelican, Harmondsworth.
 Williams, M.: 1992, *Unnatural Doubts*, Blackwell, Oxford (1998, Princeton University Press, Princeton).
 Williams, M.: 2001, *Problems of Knowledge*, Oxford University Press, Oxford.
 Wittgenstein, L.: 1969, *On Certainty*, Harper, New York.

Department of Philosophy
 The Johns Hopkins University
 347 Gilman Hall
 3400 North Charles Street
 Baltimore, MD 21218
 USA
 E-mail: mwilliams@jhu.edu

THOMAS GRUNDMANN

INFERENTIAL CONTEXTUALISM, EPISTEMOLOGICAL
REALISM AND SCEPTICISM: COMMENTS ON WILLIAMS

ABSTRACT. In this paper I will discuss Michael Williams's inferential contextualism – a position that must be carefully distinguished from the currently more fashionable attributer contextualism. I will argue that Williams's contextualism is not stable, though it avoids some of the shortcomings of simple inferential contextualism. In particular, his criticism of epistemological realism cannot be supported on the basis of his own account. I will also argue that we need not give up epistemological realism in order to provide a successful diagnosis of scepticism.

The core of Michael Williams's inferential contextualism¹ is the so-called “default and challenge” model of justification – and also of knowledge, insofar as Williams attributes a justificational component to knowledge.² On this model, a person is entitled to hold a given belief (*p*) as long as there are no justified challenges to it. If any justified challenges to the truth or reliable formation of the belief arise, the person has to cite evidence to rebut them. Otherwise she loses her entitlement to continue holding the belief. There are some propositions that entertain a default position, but are unassailable within a given context of inquiry. These are constitutive presuppositions of the respective inquiry. Williams also calls them “methodological necessities” (Williams, 2001, p. 160). When studying history, for example, one has to presuppose it is neither the case that the earth came into existence only five minutes ago nor that all historical documents are brilliant counterfeits. If one calls these presuppositions into question, then the direction of inquiry changes. One is no longer studying history but doing epistemology. Williams assumes that the methodological necessities of a given inquiry are so specific that, taken together, they determine the direction of inquiry.

If entitlement by default position depends only on what kind of inquiry we are engaged in and what challenges currently confront us, then we have the following picture: what requires an inferential justification and what does not (and is thereby basic) is not



context-independent and invariant but, rather, varies with the direction of inquiry and with the justified challenges that confront us. There is therefore no context-independent inferential structure of justification. An example will help to clarify that. An experimental psychologist supports his theories of the subjective sensations of test subjects with statements about observable behavior. So statements about the external world are basic in this context, whereas statements about subjective sensations are inferentially justified. Even in the context of psychological inquiry the basic status of individual statements about observable behavior can be undermined if there is reason to assume that an error has slipped in. In the context of sceptical considerations, on the other hand, it is a matter of course that subjective experiences have epistemic priority over beliefs about the external world. Thus the structure of inferential justification changes with the context.

If this were all Williams had to say about the conditions of justification, the result would be a justificational relativism. Another example will make that clear. In the context of astrology there is much to criticize, but not the methodological necessity that the course of the stars influences our lives. As soon as this presupposition is questioned one has abandoned the context of astrology. Of course, in the context of scientific research this presupposition can indeed be criticized. So if someone bases his predictions in the context of scientific research on the course of the stars, he does not have a justified belief, because he cannot dispel standing objections and thereby violates the conditions of the default and challenge model. But as long as someone remains within the context of astrology, these objections do not pertain to him, because they are inadmissible in this context. We cannot even say the scientist is right if he says that the astrologist holds unjustified beliefs, since Williams does not espouse attributer contextualism. For him, only the context of the epistemic subject is decisive.³ We have to accept that the astrologist holds justified beliefs in his context, but that the same evidence would not be sufficient for justified belief in science. And this is true in every context.

Fortunately, Williams wants to avoid these implications. He explicitly states that his version of contextualism does not commit one “to think of contexts of justification as insulated from external criticism” (Williams, 2001, p. 227). But how can Williams shield himself from the relativist implications? I think his externalist conditions of justification do most of the work. For beliefs to be justified in a given context of inquiry, it is not sufficient for them to be

dialectically unchallenged and for standing objections to be rebutted to satisfaction (this would be required by the weak reading of the default and challenge conception). Rather, the beliefs have to be reliably formed in an objective sense (Williams, 2001, p. 149). For this it is necessary that the methodological necessities of an inquiry are true (I will call this the strong reading of the default and challenge conception). If the course of the stars does not influence our lives, then astrological predictions are not reliable. So even if methodological necessities do not need to be *justified* within a given context of inquiry, they do have to be *true*. And this is exactly the point where external criticism can be applied. In the context of scientific research, one can arrive at the justified view that astrology is on the whole unreliable and therefore does not produce justified beliefs, since the course of the stars does not influence our lives.

Thus it is the externalist element of Williams's position that blocks justificational relativism. But this externalism has implications that I do not think are compatible with another of Williams's views – namely his epistemological anti-realism. Allow me to illustrate this with another example. Let's look at a group of particle physicists. They will only arrive at justified beliefs if they satisfy an array of methodological necessities in order to meet the externalist requirements for reliability of results. For example, the researchers must be able to identify, use and read their instruments. These and other requisite cognitive abilities are no longer at issue in the context of the specific inquiry. But they have to be present in order for reliable results and thereby justified beliefs to be produced.

But the reliability of these specific cognitive abilities is not a primitive fact. It depends on the reliability of more fundamental cognitive abilities. The researchers could not identify, use and reliably read their instruments *if they did not have reliable sense perception*. So justification in the context of specific inquiries depends on the reliability of specific cognitive abilities, which in turn depend on the reliability of abilities of some general type (such as sense perception, memory, and so on), which are generic sources of justification in many different contexts of inquiry. These generic sources need not be constituted by one single method, but they must be constituted by methods with sufficient natural similarity, i.e. perceptual processes, memory processes and so on. In other words: concealed within the context of inquiry is an externalist fundament of inferential justification, which allows just the sort of general epistemic classification of beliefs that Williams rules out on the basis of his analysis of the

inferential structure of justification.⁴ Williams seems to have in mind the following sort of argument (see, e.g. Williams, 2001, pp. 171, 193f.):

- (1) There is no context-independent, invariant structure of inferential justification.
- (2) If there is no context-independent, invariant structure of inferential justification, then there is no invariant epistemic structure.
- (3) Epistemological realism (the assumption that beliefs fall into natural epistemological kinds) presupposes the invariance of epistemic structure.
Therefore,
- (4) Epistemological realism is false.

Premises (1) and (3) seem to me clearly to be right. (1) follows directly from the default and challenge conception of justification. Whether something is in need of justification or itself can serve as a basis for justification depends on what is questionable and what is not in a given context. (3) is plausible because only an invariant epistemic structure could provide a sufficient basis for a non-arbitrary epistemic classification of beliefs. Premise (2), however, I consider false. The antecedent can be true without the consequent being true if epistemic structure is not determined exclusively by inferential facts. Strict inferentialism is sacrificed if one supplements the default and challenge model with externalist additional conditions, i.e., if one accepts the default and challenge conception according to its strong reading. Justification would then no longer depend exclusively on the inferential status of beliefs, but also on objective facts about the reliability of cognitive abilities. And if the global dependence relations that I assume to exist do in fact exist, then they allow (below the level of inferential relations between beliefs) a classification of our beliefs by natural kinds. Beliefs would then be classified according to which basic cognitive abilities (Williams also calls them generic sources) their reliability depends on.

What role does the truth of epistemological realism play for scepticism? Williams views it as follows: if the default and challenge model is right, then the possibility of Agrippian scepticism is nipped at the bud because entitlement by default position halts the infinite regress of justification (Williams, 2001, pp. 148f.). But according to Williams there is another sceptical strategy available to the Cartesian sceptic. If he can manage by constructing sceptical hypotheses to explain how our methods of belief formation could systematically lead to error, then he will automatically have produced genuine defeaters that subvert the default justificational status of the ordinary knowledge-claims that are based upon these methods.⁵ In this sense, the Cartesian sceptic would have constructed defeaters, given the

truth of epistemological realism. If beliefs fall naturally into broad epistemological classes (i.e. beliefs about the external world, beliefs about the past, beliefs about other minds) that are unified by their generic sources, and if the truth or reliability of these classes of belief in general is called into question by Cartesian sceptical hypotheses, then there is no conceivable way one could defeat these defeaters. Defeaters must be dispelled with the help of uncontroversial evidence, and such evidence for, say, our beliefs about the external world is not available if their justification in general is suspended by a defeater. It would be different if all epistemic methods really were completely context-specific. Then there always would be some epistemic method or other not impugned by the defeater that could be used in attempting to refute a given defeater. When the reliability of a certain kind of instrument becomes questionable, then it can be evaluated and perhaps defended with the help of another kind of instrument (that delivers information about the same domain).

But wait! If epistemological realism is correct, is the danger of scepticism really as great as Williams thinks? I see at least two reasons not to think so. First of all, one has to be very careful in the formulation of epistemological realism. It allows only individuation of *kinds of knowledge*, not of *domains of knowledge*. We can take perceptual knowledge, memory knowledge and introspective knowledge to be natural epistemic classifications insofar as large classes of beliefs are justified by sense perception, memory or introspection. But that does not exclude the possibility of different kinds of knowledge about the same domain. It seems for example possible that there could be not only perceptual but also a priori knowledge of the external world (if there is any such thing as a priori knowledge). And if that is so, then the reliability of sense perception could be defended against sceptical defeaters with the help of independent a priori evidence. Thus if one formulates the position carefully, epistemological realism does not lead directly to scepticism.⁶

Secondly, Williams' view that sceptical hypotheses are genuine defeaters seems suspect to me. The mere possibility of particular deception scenarios does not amount to a defeater, no matter how realistic the possibility may be. Rather, there have to be reasons that attest to the actual reality of such scenarios.⁷ Let us assume I have a visual experience of something red and then venture the claim that there is a red object before me. Someone else objects that it is possible (in the sense of a real possibility given in our world) that the object is in fact white but illuminated by red light. As I see it, this objection is not yet sufficient for a defeater. There need to be reasons that attest to

the factual or at least probable realization of this possibility. Sceptical hypotheses are in principle afflicted with the same flaw. However well they may be formulated (through brains-in-a-vat scenarios) as life-options and however much explanatory power they may possess, they are not defeaters unless we have reason to believe that they are (at least probably) realized. And this condition cannot be satisfied by sceptical hypotheses, because they are deliberately constructed in such a way that we certainly could not have any reason to attest to their realization (since the deception proceeds *ex hypothesis* unobserved). Thus they are not genuine defeaters. So in the end epistemological realism gives a lot less cause for sceptical doubt than Michael Williams thinks.

Finally, let me give a brief sketch of two more concerns I have about Williams's version of contextualism. First: Williams is confronted with a certain dilemma regarding the status of epistemological realism. Either he assumes that epistemological realism is a methodological necessity of epistemological inquiry in general, or he assumes that it is not. In my view, both alternatives create irresolvable problems for him.

If epistemological realism had the status of a methodological necessity of epistemological inquiry in general, it would be presupposed as obvious without further justification – although it is in itself anything but obvious. It would be a blind spot in the epistemologist's eye. But this picture of epistemology seems anything but plausible to me. It is true that traditional epistemologists have accepted epistemological realism without qualification, but not as an unjustified assumption. On the contrary, they have had reasons for doing so – and good ones at that, if I am not mistaken. But even if epistemological realism were false, as Williams believes, then this would in any case be a topic of discussion *within* epistemology – as Williams's own work aptly demonstrates. Williams's version of contextualism is of course an epistemological position too. I do not see how one could contest that.⁸ Anyhow, if epistemological realism can be called into question in the context of epistemology, then it cannot be a methodological necessity of epistemology.

If, however, epistemological realism is taken not to be a methodological necessity of epistemology in general, then it remains mysterious – or at least Williams has not explained – why this fundamentally false picture of our epistemic situation is anchored so deeply in our traditional conception.

Second: according to Williams, the justification of a belief (*p*) in a specific context of inquiry, C1, demands *only* that the methodological

necessities (and externalist conditions) are true and that all defeaters relevant in the given context have been dispelled. But what happens if in the context of another inquiry, C2, an unrefuted (or even unrefutable) defeater against the methodological necessities of C1 emerges, although these methodological necessities are *de facto* true? One possible answer is that (*p*) is justified by the satisfaction of the contextual conditions of C1 up until the defeater emerges in C2, but that the emergence of the defeater automatically undermines the justification of (*p*). In other words, the contextual conditions relevant to the justification of (*p*) are only provisionally valid; they are altered by the emergence of external criticism. Hence there is a change in the relevant context. If this is Williams's position, then it amounts to an epistemological externalism that admits all kinds of defeaters (without any contextual constraints).⁹ But then it would no longer be clear why this position should be labelled "contextualism". Simply accepting defeaters does not make it contextualist. Or maybe Williams sees the situation as follows: in the context of C1 we maintain our justification to believe (*p*) in the face of external criticism, although we are justified in the context of C2 to believe that (*p*) is not justified. And this is unambiguously true for any context of ascription or self-ascription since Williams does not subscribe to the view that the semantics of epistemic terms is determined by the context of the ascriber. So not only epistemic meta-inconsistency¹⁰ may occur between different contexts, it also may be recognized as such by the epistemic agent herself. If we accept this, however, we are confronted with the question why we should yield to external criticism and not just persist in holding the belief that has been criticized. Does this fragmentation of reason not contradict our intuition that justification should make rational orientation possible? Thus, Williams must decide whether he wants to give up contextualism or hang on to it at the price of a partial abandonment of the demands of rationality.

NOTES

¹ What I – along with Pritchard, 2002 – label here as "inferential contextualism" Williams in his contribution to this volume calls "issue contextualism" for the simple reason that on his view epistemic requirements change with the subject-matter of inquiry.

² See Williams, 2001, p. 149. He ascribes this term to Robert Brandom. For Williams's claim that knowledge requires justification compare *ibid.*, p. 35.

³ This follows from Williams's outline of contextually relevant factors in his 2001, pp. 159–162. For a similar evaluation, see Pritchard 2002.

⁴ For a more comprehensive defense of epistemological realism, see Grundmann, 2001.

- ⁵ In Williams, 2001, Williams makes various remarks that attest to his conviction that sceptical hypotheses can have the status of genuine defeaters. Cf. Williams 2001, p. 169: "If a challenger implies that we might be making a mistake, we are entitled to ask how. If a challenger has nothing to say – if his challenge is genuinely naked – then no real challenge has been entered. We have no idea what sort of defence is being demanded of us. (...) The sceptic's challenge is not altogether naked." Cf. also Williams, 2001, p. 186: "The Cartesian sceptic does not assume the right to enter naked challenges. Instead, he presents certain carefully chosen defeaters of ordinary knowledge-claims: sceptical hypotheses, like that of the brain in a vat. He doesn't say that we might be wrong somehow or other: he explains how we might be systematically deceived. Once entered, challenges based on these defeaters deprive ordinary knowledge-claims of default justificational status."
- ⁶ For a more comprehensive discussion of this point see Grundmann, 2003, pp. 359ff.
- ⁷ Interestingly enough, Williams himself seems to be committed to this view. See Williams, 2001, p. 149: "Appropriate defeaters cite reasonable and relevant error-possibilities. There are two main types. Non-epistemic defeaters *cite evidence* that one's assertion is false (...). Epistemic defeaters *give ground* for suspecting that one's belief was acquired in an unreliable or irresponsible way." My italics.
- ⁸ This is reason enough – as an anonymous referee of this article pointed out to me – not to assume that Williams considers epistemological realism a methodological necessity for epistemology in general. In the final chapter of his 2001, Williams sketches a picture of epistemology without epistemological realism.
- ⁹ A comparable position is espoused in Goldman, 1979.
- ¹⁰ I speak of epistemic meta-inconsistency iff *S* is justified to believe that *p* and simultaneously justified to believe that *p* is not justified.

REFERENCES

- Goldman, A.: 1979, 'What Is Justified Belief?', in G. Pappas (ed.), *Justification and Knowledge*, Kluwer, Dordrecht, 1–23.
- Grundmann, T.: 2001, 'Eine psychologische Verteidigung des erkenntnistheoretischen Realismus', in T. Grundmann (ed.), *Erkenntnistheorie*, mentis, Paderborn, 188–209.
- Grundmann, T.: 2003, *Der Wahrheit auf der Spur. Eine Verteidigung des erkenntnistheoretischen Externalismus*, mentis, Paderborn.
- Pritchard, D.: 2002, 'Two Forms of Epistemological Contextualism', *Grazer Philosophische Studien* 64, 19–55.
- Williams, M.: 2001, *Problems of Knowledge*, Oxford University Press, Oxford.
- Williams, M.: 2004, 'Knowledge, Reflection and Sceptical Hypotheses', *Erkenntnis*, 61, 313–343.

Philosophisches Seminar
 Universität zu Köln
 Albertus-Magnus-Platz
 50923 Köln
 Germany
 E-mail: thomas.grundmann@uni-koeln.de

ANTONIA BARKE

EPISTEMIC CONTEXTUALISM

ABSTRACT. Any contextualist approach to knowledge has to provide a plausible definition of the concept of context and spell out the mechanisms of context changes. Since it is the dynamics of context change that carry the main weight of the contextualist position, not every mechanism will be capable of filling that role. In particular, I argue that one class of mechanisms that is most popularly held to account for context changes, namely those that arise out of shifts of conversational parameters in discourses involving knowledge claims, are not suited to the job because they cannot account for the genuinely epistemic nature of the context shift. A form of epistemic contextualism that defines the context through the structure of our epistemic projects is suggested. Context changes in this account are linked to changes in the background assumptions operative in our epistemic projects and the methods used to carry out our inquiries.

1. INTRODUCTION

According to the main thesis of contextualism each knowledge claim has its own context, which we have to take into consideration when judging whether someone knows that p or does not know that p . According to contextualism, someone may know that p relative to context C_1 while not knowing that p relative to a different context C_2 .

In spelling out his or her position, there are two main tasks any contextualist faces:

- (1) He needs to provide a plausible explication of what constitutes a context.
- (2) He has to offer an account of the dynamics of context changes. When do contexts change and what induces the changes?

I shall begin by sketching the answers contextualists have given to these two questions and pointing out some shortcomings of these accounts before trying to supply what I think might be better answers to these questions.

2. WHAT IS A CONTEXT?

There are two basic suggestions of how to explicate the concept of a context: the first with the help of relevant alternatives, the second by means of differing standards of knowledge.



Erkenntnis **61**: 353–373, 2004.

© 2004 Kluwer Academic Publishers. Printed in the Netherlands.

If we think of contextualism as using the resources of a relevant alternatives approach, one way of spelling out the contexts C_1 , C_2 , etc. would be to regard each of them as a specific segmentation of the alternatives into those which are relevant and those which are irrelevant.

In Dretske's classical example (Dretske, 1970), when the father and the son visit the zoo, and the son asks his father which animals the striped horse-shaped creatures in the enclosure are, the alternative that they are painted mules is irrelevant and thus does not have to be excluded before the father can claim the knowledge that they are zebras. However, if the situation had been that the father is there in his capacity as a zoo inspector and is asked to report to his superior which animals are in the enclosure, the alternative that the animals are painted mules may be a relevant alternative that the zoo inspector has to be able to exclude before he can claim knowledge. For the proponent of a relevant alternatives account of knowledge, a context can be said to be the segmentation of alternatives into those that are relevant to a knowledge claim in a given context and those that are not relevant.¹

Often the contextualist position is formulated in a slightly different way by saying that what constitutes a context are the *standards* of knowledge operative in that context. When Mr Smith and his son enjoy their day out at the local zoo and the son asks what animals these are, the standards are low. According to those low standards Mr Smith knows that the animals are zebras. However, when the denial of the possible defeater "these animals are painted mules" is mentioned, the standards are raised and Mr Smith has to consider the statements relative to these new, higher standards. According to those higher standards, he does not know that the animals are not painted mules and neither does he know that they are zebras.²

This second formulation mainly comes into play when the target is scepticism. Formulated like this, the contextualism of standards is able (or so it is claimed by its proponents – a claim that has not gone unchallenged) to explain the appeal and persuasive power of the sceptical arguments on the one hand as well as their failure to induce a lasting change in the way we view our everyday knowledge claims on the other. According to the standard contextualist both the persuasiveness of the sceptical arguments and the feeling that they do not change our everyday knowledge claims have their roots in the contextualist nature of knowledge claims. If knowledge claims are intrinsically contextualistic and a mechanism operates by which the mentioning of sceptical possibilities raises the standards, the contextualist can explain why – as long as our normal standards are the

yardstick – we know most of the things which we normally think we know. The reason is that relative to ordinary contexts and the low standards associated with them we *do* know these things. The explanation of why the sceptical arguments strike us as so forceful when we are confronted with them is just as straightforward. They are so powerful because they raise the standards relative to which we have to evaluate our knowledge claims. And relative to those high standards we *do not* know even the most ordinary things. However, as soon as the standard is lowered again, the original knowledge claims are re-instated. This means that scepticism seems irrefutable as long as the sceptical standards are operative, but as soon as the standards revert to the ordinary low level, the knowledge returns. Thus, the effect of sceptical arguments is not lasting.

One can interpret the relevant alternatives approach in such a way that it is more or less equivalent to standard contextualism by spelling out what is meant by “higher” or “lower” standards in terms of relevant alternatives. Lower standards are standards according to which fewer possibilities are relevant and thus need ruling out. In ordinary situations standards are low and only few and mundane possibilities are relevant. “Higher standards” means that more possibilities are included in the set of relevant alternatives. In the case of the very high standards induced by sceptical arguments, the sceptical hypotheses are regarded as relevant alternatives and have to be ruled out. Since the sceptical hypotheses are constructed in such a way that they cannot be ruled out by empirical means, the sceptical arguments seem irrefutable. However, we should note that though similar, the positions are not equivalent.³

Moreover, speaking of standards evokes a somewhat different picture from speaking about the ruling out of alternatives. The former suggests that there is a yardstick that may vary according to the situation, which we can meet or fail to meet. The RA account, on the other hand, tends to focus more on the content of the knowledge claim and on ways in which the world could have been different. The relevant alternatives delimit the field of the things that are seriously considered as rivals to the knowledge claim under consideration.

3. WHEN DO CONTEXTS CHANGE AND WHAT INDUCES THE CHANGE?

Let us now turn to the second point, how to account for the dynamics of context changes. Most authors working in this area have such an account. In their view the changes in context arise out of changes of

conversational context. It is the dynamics of conversation that induce context change. Let us call this position *conversational contextualism*.⁴ I shall later contrast conversational contextualism with a form of contextualism that one might call *epistemic contextualism*.

According to conversational contextualism, the way in which we can change what counts as relevant, is by mentioning it or drawing attention to it in some other way. By mentioning the possibility that the animals are not painted mules, the possibility that they are becomes relevant. Through this the context has changed and now includes this possibility. The changed context then requires a fresh evaluation of what is and what is not known.

Famously, David Lewis has proposed an account along these lines. (Lewis, 1979, 1996) Why should merely mentioning something make it relevant? Lewis offers an instance of a more general conversational rule, the so-called rule of accommodation, which says that if something is required for an utterance in a conversation to make sense, then, all else being equal, this sets the score of the conversation at this value. The rule of accommodation governs conversational contexts in general, not just conversations about knowledge. It ensures that upon an utterance being made the conversational context adjusts in response to what is said so as to make it acceptable (not necessarily true, of course, but at least capable of being true). The relevant alternatives includes adjusting some contextual parameters such as salience, presuppositions and standards.

There are several problems using the rule of accommodation to account for epistemic context changes:

Firstly, since the dynamic is a conversational one, it cannot explain why these changes happen when no conversational context is involved, as in solitary reflections on some issue. Even if Mr Smith only considers “these are zebras” and the entailment “these are not painted mules” all by himself, there would still be a question whether he knows the second statement and this could not be explained as a change in *conversational* context. In response to this, one could of course hold that solitary considerations of questions like these are conversations with oneself. However, this does not seem to be a very attractive solution.

Secondly, conversational accounts have to explain the asymmetry of why we can raise the standards for knowledge conversationally, but not lower them. For other context relative terms such as flatness we are able to do both by means of conversational mechanisms. Why, in the case of knowledge, is it so much easier to raise the standards than to lower them?

Thirdly, the conversational contextualist has to explain why, despite this, after a while the standards revert downwards of their own accord. Consider again the zebra case. It seems we can raise the standard so that the painted mules possibility becomes relevant by being mentioned. Interestingly, this raises the standards regardless of its formulation, even if it denies the sceptical possibility: “Ok, so then I also know that these animals are not painted mules.” However, we cannot lower the standards again so easily. A return to the original standard cannot be achieved simply by mentioning the original claim again. In the case of an accommodation of presuppositions that have to be made to render statements acceptable in terms of reference, we can easily shift back and forth. Also in the case of the standards of flatness we can easily return to the lower standards by conversationally returning to the discussion of roads rather than pool tables. So why in the case of accommodating standards for knowledge do we find the shift towards stricter standards or the inclusion of more relevant alternatives easier than the lowering of standards or the exclusion of possibilities?

In my view, taken together the second and the third point strongly indicate that there is more to the changes in standards for knowledge than simple, revocable conversational shifts as in the cases of reference, salience or standards like those applying to flatness etc. As we shall see the reasons why we cannot shift back as easily are *epistemic* rather than conversational.

A fourth problem arises from a conversational account: an account that makes the mentioning of an alternative so potent renders impossible discourse about what is and what is not relevant in any given case (and the potential use of such discourse should be obvious).

This last feature has been spelled out very clearly by David Lewis (1996) although it is also inherent in the other conversational accounts.

Knowledge according to Lewis (1996)

S knows that *P* iff *S*'s evidence eliminates every possibility in which not-*P* [...] except for those possibilities that conflict with our proper presuppositions.

How convincing this definition is going to be obviously depends on the way in which one fills the expression “proper presuppositions” with content. This terminology translates smoothly into the relevant alternative way of speaking. What we properly ignore are the

irrelevant alternatives: but which alternatives are irrelevant to a knowledge claim? Lewis's answer consists in a list of seven rules, which specify when we can ignore something and when we cannot. He concludes the enumeration of the rules with a rule that he himself calls "more a triviality than a rule," the so-called "rule of attention."⁵ This rule states that "ignoring a possibility" means actually ignoring it, not that one could have ignored the possibility in question, but did not. This means that a possibility can only be ignored if it is not mentioned and attention is not drawn to it in any way. As soon as attention is drawn to a possibility, it is not ignored any longer and *a fortiori* is not ignored properly.

At first sight, this seems to explain what puzzles us about examples such as the zebra case. In the zebra case one mentions the painted mule possibility when considering the entailment. By doing this one has turned one's attention to this possibility and cannot legitimately ignore it thereafter.

Given that the rule of attention has to carry the burden of this account, we should have a closer look at it. Lewis's rule, which just makes explicit a feature that is implicit in all conversational contextualist accounts, has the implausible result that we do not have any way of communicating about our presuppositions in a manner that is not pragmatically self-defeating. As soon as one of the participants of some conversation mentions a possibility, even by mistake, one cannot ignore it anymore, although everyone ignored it properly before. The only way in which one can dispose of this possibility again is by simply beginning to ignore it again. In order to do this, we cannot agree on ignoring it henceforth, but rather everyone has to divert their attention away from it. Lewis even thinks it may be necessary to break off the conversation and restart it at a later time, to have a "fresh start" in ignoring.

Surely, this cannot be right. We could just discuss the matter and then agree that it was a mistake. At least, *if* we think that it can at all be a route to knowledge to ignore the possibility in question, why should we be forced to reach this agreement silently and proceed without mentioning the presupposition again?

Moreover, discussing our presuppositions and *agreeing* on the question of which ones are reasonable (i.e. which alternatives we can reasonably ignore) and which ones are not is an integral part of our epistemic practices. According to Lewis and other conversational contextualists, it could not be. Lewis's view implies that it is not only scepticism, but also epistemology that threatens our ordinary knowledge by attending to its presuppositions. Epistemology thereby

sweepingly achieves what mentioning the painted mule possibility did in the zebra case – it brings to our attention hitherto ignored (and mainly properly ignored) possibilities and robs us of the ignorance that protects our knowledge as long as we do not engage in epistemology. In my view, this is not an attractive position.

Accepting the rule of attention means that the epistemic admissibility of our presuppositions is directly dependent upon what contingently enters our mind. Another way of formulating the problem with Lewis's account, therefore, is to say that the rule of attention conflates psychological parameters with epistemic ones: the psychological level (what we can or have ignored *de facto* in a psychological and contingent sense) is confounded with the epistemological level (what we may ignore, epistemically speaking). What is "proper" or not, epistemically speaking, must be independent of what contingently enters my mind.

To sum up: conversational rules are the wrong type of rule to govern the context changes we are interested in because conversational rules can only show us what is *conversationally proper*, but not what is *epistemically adequate or rational*. I take it that these arguments show that the conversational model of context change is not capable of providing an epistemically convincing core to contextualism.

4. EPISTEMIC CONTEXTUALISM

The discussion of conversational contextualism leads to the recognition that what we should really look for is a mechanism that could explain the specifically *epistemic* dynamics of context changes. We can extract four features a satisfactory type of contextualism should have.⁶

- (A) A really satisfactory account would be able to point to a mechanism that is epistemic in nature and thus able to make the dynamics epistemically intelligible.
- (B) The contextualist account should be capable of explaining the well-known asymmetry that it is easier to raise the standards for knowledge than to lower them again. With regard to other context-relative terms such as "flat" or "empty" standards can be raised and lowered with equal ease by conversational means, but this is not so with regard to the standards of knowledge. Conversational contextualism finds it difficult to account for this difference.
- (C) An important intuition people usually have when confronted with the zebra case is that a kind of circularity is present.

However, the accounts we have reviewed so far are silent on this circularity. Ideally, our account would explain that feeling.

- (D) Some contextualist accounts have difficulty explaining cases in which we gain knowledge by deductive inference. If, for example, mentioning the consequence always changes the context, then how can deduction be a route to knowledge? As soon as we mention whatever we deduced, we run the risk of changing the context and thereby losing the knowledge we have just gained by deducing it. If one does not wish to subscribe to the view that deduction never yields knowledge, a contextualist account has to provide criteria that distinguish between cases in which we *can* achieve knowledge of the conclusion by means of deduction from cases in which we *cannot* do so.

Desired features of a contextualist account:

- (A) Context changes should be epistemically motivated.
- (B) The intuition of circularity should be explained.
- (C) The asymmetry regarding raising and lowering of standards should be explained.
- (D) There should be clear epistemic criteria which differentiate between cases in which deduction is knowledge-yielding and those in which it is not.

In order to develop an account that meets these conditions, I shall start with some brief remarks about our epistemic projects and their general structure. I shall then combine this with the main insight of contextualism. This will yield a contextualist account that is an *epistemic* rather than a conversational contextualism.

Belief Systems

At any point in time, we have many beliefs about a variety of things, while we do not have any beliefs regarding other matters. Our beliefs stand – among others – in logical relations and epistemic relations. Our belief system does not have to be logically transparent to us – there are logical relations among my beliefs that are too complex for me to recognize and others that I *de facto* have not noticed.⁷

Questions

An obvious but interesting feature of human belief systems is that they contain gaps. Some of those gaps I might want to fill. They

are my *questions*. Of course, not every gap gives rise to a question. There are many things about which I do not have any beliefs and am not interested in having any. For instance, I do not have any belief about the exact number of hairs on my head, but I am not particularly interested in acquiring such a belief, either. Questions, in my view, are central to the dynamics of our epistemological endeavors. Among other obvious uses of questions a central function of questions is to generate and structure our inquiries.⁸

Inquiries

Suppose we have a question. Now we are interested in the answer to our question. How do we arrive at an answer? We cannot expect answers to our questions simply to pop into our heads – we have to *do* something to find the answers. In order to find the answer to our question, we conceive of an *inquiry*, which, we hope, will provide the answer.⁹ Let us define an inquiry by the question and the method used to answer it. Therefore, my inquiries are in two ways dependent on my other beliefs. Firstly, since my questions arise out of gaps in my belief system, they are dependent on what other things I believe. Secondly, they are dependent on my beliefs because in conceiving of an inquiry, I have to choose a method that I think will yield an answer to my question. Thus, I choose a method for the inquiry on the basis of the other things I believe.

Methods

The concept of a method is central to our inquiries and will play an important role in the account I shall seek to develop, so I shall say a few brief words about some basic features of methods. In general, methods are procedures that are aimed at the realization of a more or less specific goal. In order for a procedure to be a method, it has to enable us – at least in principle – to achieve the goal in question by use of the procedure. This implies that the procedure may only have a finite number of steps and that it has to lead to the achievement of the goal in a reasonable number of cases. In order for any given method to work, initial conditions and framework conditions have to be fulfilled. A good method of lighting a fire is to rub a match along a match box and to bring the burning match close to a pile of paper. This is a good method because it leads to success in most instances. However, it will only lead to success if certain initial conditions and framework conditions hold, such as

that the match, the strip and the pile of paper are dry, oxygen is present, the natural laws hold, etc.

“Epistemic methods” are a particular sub-set of our methods, namely those that serve epistemic goals. Epistemic methods, therefore, differ from other methods mainly with regard to their aim: the different epistemic methods are directed at finding the answers to a variety of questions.

Epistemic Projects

Some of our epistemic projects are highly sophisticated, meticulously planned enterprises, such as the search for the cause, or causes, of schizophrenia; some are much more mundane, such as finding the answer to the question whether there is still cheese in the fridge. Some do not really feel like “projects” at all, such as finding out what is right in front of my eyes. There are epistemic projects that I can carry out on my own, whereas others can only be done as a co-operative enterprise together with other epistemic agents. My epistemic projects can be short and of interest only at a particular point in time (the cheese project again), or they may take considerable time and the result may be important for even longer, such as the schizophrenia project. Finally, and importantly, epistemic projects are themselves structured and may consist of many individual inquiries. For instance, the project of finding the causes of schizophrenia itself consists of many individual questions that have to be answered and the related inquiries.

Assumptions

With these concepts in place, I can now say more about the concept of assumption that plays a central role in the account I shall develop. I have already mentioned in connection with methods in general that using a method imposes certain requirements. Let us now consider these requirements with respect to epistemic methods. Like all methods, epistemic methods can only be successful if certain conditions are *in fact* fulfilled. In order for me to be able to determine the temperature of some water in a beaker by using a thermometer, the thermometer has to be functioning. In order for me to find out when the next train leaves, the timetable has to be genuine. In order for me to find out the colour of an object through simply looking at it, my eyes have to be in good condition and the lighting relatively normal, and so on. Over and above that, when I *deploy* a method, I have to *assume* (for its use) *that these conditions are fulfilled*. This does not mean that “assuming that conditions $c_1, c_2, \dots c_n$ are fulfilled” is a conscious

mental activity; it only requires that I take these conditions – defeasibly – to be fulfilled. For most of them, I can of course choose to investigate whether they are fulfilled (and sometimes I might have to do this). But such an investigation is another inquiry, which requires different methods and different assumptions. I can use a thermometer to measure the temperature of some water. In this inquiry I assume that the thermometer is in working order. Of course, I can *test* whether the thermometer is in working order. However, testing whether a thermometer works properly is a different inquiry from measuring the temperature of some water. So I have generated a *new inquiry*, which seeks the answer to a different question, namely “is the thermometer in working order?” rather than “what is the temperature of the water?”. This new inquiry requires a different method: I could for example use the method “comparison with a standard thermometer”. Obviously, this new method also requires the fulfilment of certain initial and framework conditions. The standard thermometer has to work properly etc. That these conditions are fulfilled has to be assumed with regard to this new inquiry. It is interesting that another method of testing the functioning of the thermometer would be to dip it into water of a known temperature and to ascertain whether it gives the correct reading. This is an important point because it shows that the very same physical event (and the very same physical laws etc.), in this instance immersing the thermometer in the water and reading off the measurement, can be part of different methods, depending upon its status in our project, that is to say on the question asked and the assumptions made. The interpretation of the measurement read off is not determined by anything intrinsic to it, but by the status it has within the framework of the inquiry and its assumptions.

It is important to distinguish these two levels: in order for the method to *work* it is necessary that the conditions are *de facto* fulfilled. However, in order for someone to be able to *deploy* the method, the person using it has to *assume* that they are fulfilled (if she has not investigated it beforehand, in which case this principle would apply to her earlier investigation). For me to be able to use a thermometer in the way I do, I have to assume that certain facts hold for it. If I did not, I could not interpret the results as I do. These assumptions, naturally, are not indefeasible.

There are some types of assumption that have to be made in every investigation.

In each inquiry it has to be assumed that ...

- (a) The chosen method *M* is an adequate method for the investigation of *p*.

In particular, in order to be a method for finding out whether *A* or *B*, it has to be sensitive with regard to *A* and *B*, i.e. it has to discriminate between them. This incorporates intuitions of DeRose (1995) and Goldman (1976).

- (b) The method *M* has been applied correctly. That is to say, one has not committed any mistakes in its application, nor were there any problems with instruments, sense-organs etc. They are all (defeasibly) assumed to be in working order.
- (c) The requirements on the side of the world are fulfilled for method *M* to be applicable. The natural laws are stable etc.

The assumptions that have to be made for a given method range from the very general to the very specific. Among the most general assumptions is that there is an external world, that this world is structured in a lawful manner and that those laws are stable (i.e. they do not fluctuate from moment to moment). Among the more specific assumptions are those that have to be assumed for e.g. particular types of measurement. Suppose you want to find out the buoyancy of an object by weighing the liquid it displaces. For this, you have to assume that the gravitational constant $g = 9.81 \text{ m/s}^2$. An even more specific assumption for the use of this method would be that the particular scales you use are in working order.

Whether a method counts as adequate is not arbitrary, but may depend on several factors, some of them of a more epistemic nature, some of them more pragmatic. Among the epistemic factors I would count the following: Is the chosen method a method for finding out whether *p* at all? Suppose, I want to measure the weight of an object, then taking its length with a ruler would not be an adequate method. Provided the method is a method in the sense that it does lead to a result regarding *p*, does it discriminate between those possibilities that are relevant to the person's question? In most instances consulting someone's birth certificate is a method of finding out who the person's father is. However, if a man doubts whether his child is really his own, consulting the birth certificate is not an adequate method of finding out; rather, he would have to embark on a paternity test.

Among the more pragmatic factors that determine whether a method is adequate, may be how much depends upon arriving at the right result. In order to measure the temperature of my bath it would be an adequate method to immerse a household thermometer and read off the temperature. However, if I am measuring the temperature of water serving as a reactor coolant, this method should probably be judged inadequate. One could argue that similarly there are higher demands on methods of an expert than of a lay person. For a lay person it may be a wholly adequate method of determining the cause of death of a person if he saw the person being hit by a falling branch. A coroner in answering the same question may have to use different or additional methods.

Questioning Assumptions

Assumptions may be made as long as nothing upsets them or calls them into question. However, *if* anything happens that upsets my assumptions or calls them into question, I have to react to it or else I jeopardize my inquiry. In this situation, three basic options are available.

- (1) I could argue that the assumption is legitimate or plausible and does not need investigation. This option is only available if what has called the assumption into question is not evidence against it. If we already have indications of the assumption's falsehood, a defence of its epistemic legitimacy cannot be given. It cannot be legitimate to assume something if we have evidence to the contrary.
- (2) I could choose another method of finding the answer to my original question, namely a method that does not rely on the assumption that has been called into question – provided such an alternative method is available.
- (3) I could choose to investigate whether the assumption holds.

Reactions to the challenge of an assumption

- (1) Argue for the plausibility or legitimacy of the assumption.
- (2) Choose a method to investigate the question that does not rely on the assumption.
- (3) Investigate whether the assumption holds
 - (a) if it holds, proceed with original inquiry;
 - (b) if it does not hold, go to (2), i.e. choose a method that does not rely on the assumption.

Inquiries as Mechanisms of Context Change

Let us now combine these thoughts with the main insight of contextualism, namely that knowledge is knowledge relative to a certain context. Earlier, we demanded that the mechanisms that regulate the contexts should not be conversational, but epistemic mechanisms. We are now in a position to suggest such a mechanism: in my view the epistemic mechanism behind context changes should be sought in the features of our inquiries that I have just outlined. We may define a context through the inquiry and the assumptions made in order to carry it out. Context shifts are not shifts in any dimension external to the epistemic realm such as conversational parameters, but certain shifts in the assumptions and thus in the inquiry.

Let us begin by considering a single inquiry before turning to the constellations that are possible with regard to combinations of inquiries. Imagine a doctor who wants to find out about your blood pressure. The doctor has all sorts of beliefs about you, about blood pressure and ways of measuring it. She does not have a belief about your blood pressure at this moment. However, she is interested in acquiring such a belief. She asks the question ‘what is your current blood pressure?’ and frames an inquiry to find the answer. In order to do so, she has to assume quite a few things, among them some physiological theory, certain facts about the sphygmomanometer (it is in working order etc.), about the procedure (it is being done correctly, the inflatable cuff has the correct size for your arm) as well as about her own sense organs, the natural laws and so forth.

Larger epistemic projects may consist of smaller ones. For instance the question about your cardio-pulmonary health would consist of questions, say, about your ECG, your blood pressure and your level of blood cholesterol. These more limited inquiries may be combined as parts of a larger project. If more than one inquiry forms part of one larger project, certain constraints have to be kept in mind. We have to take care that the assumptions of the individual inquiries that are combined in the whole project are not incompatible, or else we lack a coherent method to carry out the whole project despite the fact that we may have methods for individual stages of it. Similarly, if we engage in co-operative projects, we have to co-ordinate our methods and assumptions. Put in general terms, we can combine two inquiries A_1 and A_2 into a larger inquiry B , provided there is no incompatibility in the assumptions that are required for the investigations A_1 and A_2 . If the assumptions we have to make for A_1 and A_2 are not compatible, such a combination is impossible. Systematically,

one way of expressing this would be to say that we would *lack a coherent method* for the investigation of the overarching question. We have said that in order for a method to work, certain conditions need to be fulfilled and in order for us to be able to deploy the method we have to assume that they are fulfilled. Suppose that among the conditions that have to be fulfilled for method M_1 there is the condition c_1 “the atmospheric pressure is normal” while for method M_2 among the conditions that have to be fulfilled there is the condition c_2 “the atmospheric pressure is below normal”. We could not combine the methods M_1 and M_2 into one method because we would at the same time have to assume c_1 and c_2 . In other words we would have to assume that the atmospheric pressure is normal and below normal at the same time. Thus we could not formulate a method M_1 and M_2 because it would require the assumption of simultaneously normal and below normal atmospheric pressure. Another way of formulating this point would be to remember that our assumptions are defeasible. That is to say, we may only assume something as long as nothing calls it into question. Suppose we use method M_1 and in order to do so we (defeasibly) assume c_1 . Now we use M_2 and in order to do so, we have to assume c_2 , a contrary to c_1 . Although we may not know *which* assumption is false, we can know that they cannot both be true. Thus, the assumption c_1 and c_2 is defeated and we may not assume it.

As long as the total set of our assumptions is consistent, we can combine the assumptions and consider the combined steps as one overall method of answering the overarching question. In this case, the context would be continuous between the projects. There is no upper limit to how long the context of a single project may last or how large the project may become. As long as we do not call into question our assumptions, the context remains stable. I have labored this point somewhat because many of the conversational contextualist accounts make it difficult to see how we could achieve integration of our knowledge into larger contexts – something that seems eminently desirable to me.

Context changes

A change in context normally takes place when we call into question one or more of the assumptions of an inquiry. Once this happens, we have to suspend our inquiry until we have addressed the problem, because until it has been settled whether the assumption in question is true or legitimate, it may no longer be assumed. There is a way in which

we can *resist* the context change when one of our assumptions has been called into question. We can defend our assumption as legitimate or reasonable, that is to say as following from other things that are not under dispute or with recourse to other standard assumptions, or, in the interpersonal case, by appealing to shared assumptions (as in (1)). Only if this is successful can we resist the context change. If we do not choose this route, or if we are not successful in establishing the legitimacy of the assumption, we have to react to the challenge in one of two ways, namely either drop the assumption and change the method that depends on it, or else investigate the assumption. If we do either of these, the context changes. If we drop the assumption and use a different method to investigate our original question (as in (2)), we have created a new inquiry in order to answer the question. This means, we have accepted the new context and are tailoring our inquiry to it. If, on the other hand, we decide to investigate the assumption (as in (3)), we ask a new question, which requires that we conceive a new inquiry and thereby changes the context.

Circularity

Now we are in a position to account for the zebra case. In the original situation the question is “what animals are the animals in the pen?”. We assume that we are in a normal public zoo, i.e. a zoo that operates normally, displays genuine specimens etc. In such a zoo, the method used, namely looking at the animals under good lighting conditions and reading the plate in front of the pen, is a means of answering our question. As long as the assumptions are not called into question, Mr Smith knows that the animals are zebras. The fact that they are not painted mules (or rather the fact that in a state zoo real specimens will be on display) is *assumed* to make the method of looking at the animals an adequate method for answering the question. Using an adequate method generates a warrant that *p*. In virtue of this warrant we know that *p*, provided the conditions that we assume are fulfilled are *de facto* fulfilled. If we are now interested in the question of whether the animals are painted mules, we cannot conclude this from the fact that they are zebras as long as we have arrived at the belief that they are zebras in a way which relied on the assumption that they were not painted mules. This would be just as circular as determining the temperature of a liquid with a thermometer and then using the liquid of known temperature (known with the help of that very thermometer) to test the functioning of that same thermometer. This

is the reason why there is such a strong feeling of circularity about the zebra case. This analysis is further strengthened by the fact that *had we arrived* at the belief in a way that did *not* rely on the assumption that the animals are not painted mules, say had we arrived at “these are zebras” by means of a genetic test (a method, which itself requires certain assumptions, but different ones) there would have been no problem with gaining knowledge of “these are not painted mules” by deducing it from “these are zebras”.

What is known in this case can be described most clearly if we acknowledge the difference between being rationally committed to something and knowing it. We are rationally committed to “these are not painted mules” in the first inquiry. However, by saying that we are rationally committed to “these are not painted mules,” we are not claiming that we *know* that they are not. There are different ways of being rationally committed to something and one of them is the way in which we are committed to the assumptions of our inquiries. However, this is not linked to differing degrees of *confidence*. It was recognized in different ways by Peter Klein (1981) as well as by Keith DeRose (1995) that I should be no less confident with regard to “these are not painted mules” than I am with regard to “these are zebras.” This is true, but the reason for this is that we should be at least as confident with regard to our assumptions as we are with regard to those things we investigate on their basis.

Lowering the standards

Epistemic contextualism also explains another feature that has puzzled previous analyses. These accounts suffered from an asymmetry that was hard to explain. Conversationally, the standards could be raised (or alternatives could become relevant), but it appeared that the reversal of this process was not possible. Standards could not be lowered conversationally in the same way they could be raised. Rather, proponents of conversational contextualism had to resort to psychological mechanisms, such as forgetting the possibility that had been mentioned or allowing it to move out of the focus of one’s attention in some other way. This contrasted with other cases of context dependency, such as “flatness” for example, in which it is possible to adjust the standards *upwards or downwards* by conversational means. A convincing explanation of this asymmetry in the case of knowledge could not be given. Epistemic contextualism can explain why we cannot simply stick to the so-called “lower” standards

and resist the pull towards the “higher” standards. This is ruled out by the basic structure of our inquiry. Once the initial assumptions have been challenged, we have to react to the challenge. Simply waiting for us to forget the challenge is not possible because the methods we have chosen depend upon those assumptions. The asymmetry reflects the fact that it is easier to challenge assumptions than to answer such a challenge.

However, at this point we have to distinguish between two fundamentally different cases. The first one is a case in which *one or more* of the specific assumptions underlying a particular project are called into question: Is this thermometer in working order? Are we in a normal zoo? Is the lighting normal? The second case is that of the sceptical arguments, in which the assumptions of any empirical project whatsoever are called into question. I shall not deal with the latter case here.

In the first case we cannot – and neither do we, in the normal course of events – simply revert to our original assumptions after one or more of the assumptions have become questionable. We *may* return to them without further investigations, but only because we judge that everything considered, they are still very plausible given all our other background beliefs. However, this should *not* be regarded as a simple “reverting” to the assumptions. Rather, it normally involves some justification or defence of the assumptions even if this is not done by a new investigation. Typically, people react to the zebra case with a story of what else they believe about zoos etc. This amounts to a defence of the plausibility of the assumption that genuine specimens are on display. As long as these other beliefs are not dislodged, the assumption that the zoo displays genuine specimens is legitimate and thus looking at the animals is an adequate method for generating knowledge. If, on the other hand, we cannot make plausible our assumption, then no amount of simply ignoring such considerations will give us back the warrant for *p*. Of course, we can simply “change the topic” in a conversation, but this cannot restore the assumptions and therefore will not allow us to conclude our inquiry in an epistemically satisfying way.

Deduction

Epistemic contextualism also provides us with a clear and intuitive criterion that allows us to differentiate between cases in which deduction is knowledge-yielding and cases in which it is not. The

criterion is independent of conversational contexts. Instead it follows the structure of when we may (in keeping with our project and the status of p and q) infer q from p . Note that the concept of circularity employed here is of an epistemic kind, i.e. connected to the structures of our epistemic projects and not a purely logical relation. Deduction is knowledge-yielding in all cases in which the claim to be deduced does not have the status of an assumption with regard to the investigation that we used to achieve a warrant that p in the first place. For this reason the inference from “these are zebras” to “these are not painted mules,” though a valid inference, is knowledge-yielding only if we arrived at “these are zebras” in a way that did not rely on “these are not painted mules” or “this is a zoo that displays genuine specimens.” The inference is not knowledge-yielding in cases in which we have made one of these assumptions and arrived at “these are zebras” by way of their appearance.

Adequacy Conditions

It is now time to review how this account fares with regard to the adequacy conditions we laid down at the outset. As desirable features of an account we isolated the following characteristics:

- (A) Context changes should be epistemically motivated.
- (B) The intuition of circularity should be explained.
- (C) The asymmetry regarding raising and lowering of standards should be explained.
- (D) There should be clear epistemic criteria which differentiate between cases in which deduction is knowledge-yielding and in which it is not.

The account developed meets all four conditions. The linking of context changes to structural features of our epistemic projects ensures that the changes are motivated epistemically rather than conversationally or in any other way extrinsic to the epistemic realm. The feeling of circularity is accounted for, since it is indeed a circular procedure to deduce q from p , if q was assumed in order for us to be warranted in p the first place. Thirdly, the asymmetry between raising and lowering the standards has been explained, and fourthly, a clear criterion has been provided for separating cases in which deduction is knowledge-yielding from cases in which it is not. Deduction is knowledge-yielding as long as we do not deduce propositions that we had to assume for the inquiry that led us to be warranted in the premises.

ACKNOWLEDGMENTS

This paper is dedicated to the memory of James Logue. Dr Logue was one of the examiners of my thesis, which was a direct precursor of this paper. He died unexpectedly on 20 January 2004.

NOTES

- ¹ A variety of relevant alternatives formulations have been used by Dretske (1970), Stine (1976), Goldman (1976), Shatz (1981), Heller (1989) and White (1991) among others.
- ² Formulations using the terminology of standards are used by DeRose (1995), Lewis (1996) and Schiffer (1996) in his criticism of contextualist solutions to scepticism.
- ³ See Barke (2002, and pp. 92f).
- ⁴ Among the conversational contextualists are Stine, Lewis, Heller, and to a certain extent DeRose.
- ⁵ For a more detailed treatment of Lewis' account see (Barke, 2002, pp. 93–103).
- ⁶ In my 2002 I discuss one further condition, namely that the condition one can distill from Schiffer's criticism (1996) is fulfilled. See pp. 143–147 and 166–170 for the treatment of this condition.
- ⁷ Here I differ strongly in my terminology from authors who are concerned with the modeling of epistemic states and epistemic logic such as Hintikka (1962) and Gärdenfors (1988), who require complete logical transparency for what they name "belief systems."
- ⁸ Philosophers who have been particularly interested in questions are among others Hilpinen (1986) and Hookway (1996). White (1991) has applied the terminology of questions and concerns to the discussion of closure and knowledge and generated an account that could be classed as occupying a half-way position between conversational and epistemic contextualism.
- ⁹ For an inspiring in-depth treatment of inquiries see Hookway (1990).

REFERENCES

- Barke, A.: 2002, *The Closure of Knowledge in Context*, Mentis Verlag, Paderborn.
- DeRose, K.: 1995, 'Solving the Skeptical Problem', *Philosophical Review* **104**, 1–52.
- Dretske, F.: 1970, 'Epistemic Operators', *The Journal of Philosophy* **67**, 1007–1023.
- Gärdenfors, P.: 1988, *Knowledge in Flux. Modeling the Dynamics of Epistemic States*, MIT Press, Cambridge MA.
- Goldman, A. I.: 1976, 'Discrimination and Perceptual Knowledge', *The Journal of Philosophy* **73**, 771–791.
- Heller, M.: 1989, 'Relevant Alternatives', *Philosophical Studies* **55**, 23–40.
- Hilpinen, R.: 1986, 'The Semantics of Questions and the Theory of Inquiry', *Logique et Analyse* **29**, 523–539.
- Hintikka, J.: 1962, *Knowledge and Belief. An Introduction to the Logic of the Two Notions*, Cornell University Press, Ithaca.

- Hookway, C.: 1990, *Scepticism*, Routledge, London.
- Hookway, C.: 1996, 'Questions of Context', *Proceedings of the Aristotelian Society* 1996/1, 1–16.
- Klein, P. D.: 1981 *Certainty: A Refutation of Scepticism*, University of Minnesota Press, Minneapolis.
- 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.
- Schiffer, S.: 1996, 'Contextualist Solutions to Scepticism', *Proceedings of the Aristotelian Society* 1996/3, 317–333.
- Schatz, D.: 1981, 'Reliability and Relevant Alternatives', *Philosophical Studies* **39**, 393–408.
- Stine, G. C.: 1976, 'Scepticism, Relevant Alternatives, and Deductive Closure', *Philosophical Studies* **29**, 249–261.
- White, J. L.: 1991, 'Knowledge and Deductive Closure', *Synthese* **86**, 409–423.

Institut für Psychologie
Johann Wolfgang Goethe-Universität
Mertonstraße 17
60054 Frankfurt
Germany
E-mail: barke@em.uni-frankfurt.de

FRANK HOFMANN

WHY EPISTEMIC CONTEXTUALISM DOES NOT PROVIDE
AN ADEQUATE ACCOUNT OF KNOWLEDGE:
COMMENTS ON BARKE

ABSTRACT. According to Antonia Barke's version of contextualism, epistemic contextualism, a context is defined by a method and its associated assumptions. The subject has to make the assumption that the method is adequate or reliable and that good working conditions hold in order to arrive at knowledge by employing the method. I will criticize Barke's claim that epistemic contextualism can provide a more satisfactory explanation or motivation for context shifts than conversational contextualism (in particular, David Lewis's contextualism). Two more points of criticism will be presented, which are meant to show that epistemic contextualism presupposes epistemic internalism, and that (epistemic) contextualism leads to an implausible view about which parameters the special achievement that is constitutive of knowledge depends on. I suggest that, contra (epistemic) contextualism, knowledge is a more robust phenomenon that does not depend on whether anyone calls into question any assumptions or raises skeptical doubts in conversation or in his or her mind (as, for example, Fred Dretske's account says). I indicate how this can be reconciled with the phenomenon that knowledge attributions are somewhat unstable and seemingly context-dependent.

1. INTRODUCTION: SOME IMPORTANT FEATURES OF EPISTEMIC
CONTEXTUALISM

Antonia Barke presents a new and interesting kind of contextualism, "epistemic contextualism," as she calls it. Her point of departure is the rejection of conversational versions of contextualism, i.e., those versions of contextualism that conceive of the dynamics of context shifts as governed by conversational facts. Her criticism of conversational contextualism, and particularly of David Lewis's version of it, is convincing. The goal then for Antonia Barke is to develop a different kind of contextualism that is "epistemic rather than conversational" in character (Barke, 2004, p. 357). According to this epistemic contextualism, the dynamics of context shifts are supposed to be *epistemically motivated*, and not just a matter of brute psychological and conversational facts, and thus is supposed to avoid the



Erkenntnis 61: 375–382, 2004.

© 2004 Kluwer Academic Publishers. Printed in the Netherlands.

charge of “conflat[ing] psychological parameters with epistemic ones” (Barke, 2004, p. 359). I understand the concern and the goal, but as I will try to show below, I do not think that Barke has succeeded in reaching the goal. In particular, my second worry below will indicate why Barke’s contextualism is not much better off than conversational contextualism with respect to displaying the supposed epistemic nature of context shifts.

Another important feature of Barke’s contextualism is that the context, with respect to which knowledge is relative, is defined by a *method* and its *assumptions*. This means, most importantly, that the proper use of a method requires that the subject make certain assumptions about the method and the circumstances. The most important assumptions here are: that the method to be used is adequate or reliable, and that good working conditions hold (Barke, 2004, p. 361; cf. Barke, 2002, p. 155). These are assumptions the subject has to make in order to acquire knowledge by use of the method. Now, these assumptions can of course be challenged, and the result of such a challenge will in effect be a raising of the standards for acquiring knowledge. In Barke’s view, it will be true then that, unless the subject can react in a certain legitimate way to the challenge, she will no longer be allowed to make the assumption. It is, so to speak, a consequence of the “rules for making assumptions” that if an assumption has been challenged the subject either has to provide some *independent* reason for making it or she really has to drop it. In the end, these are the only two legitimate ways of responding to the challenge of an assumption.

2. THREE WORRIES

My first worry is that Antonia Barke seems to *presuppose epistemic internalism* (in a sufficiently strong version) when she is introducing “assumptions.” An epistemic externalist does not agree that one has to make the “assumptions” that Barke claims one has to make in order to use a method. According to epistemic externalism, it is sufficient that certain conditions in the subject’s environment actually obtain in order for the subject to acquire knowledge. It is not necessary that the subject believes that these conditions obtain. (Epistemic externalism is well supported both by intuitions and systematic arguments, as presented in the works of A. Goldman, F. Dretske, R. Nozick and others.) For example, one does not have to believe (nor does one have to be justified in believing) that sense perception is

reliable in order to arrive at knowledge by means of sense perception. In contrast, Barke claims:

In order for me to find out the colour of an object through simply looking at it, my eyes have to be in good condition and the lighting relatively normal, and so on. Over and above that, when I deploy a method, I have to assume (for its use) that these conditions are fulfilled. (Barke, 2004, p. 362; cf. Barke, 2002, p. 155)

The condition mentioned in the first sentence is fine, but the second sentence presupposes epistemic internalism. So an externalist will already have general reasons for finding Barke's approach unconvincing.

Now, perhaps, "assumptions" means something different here. But then, we would have to be told more about what it involves. And I suspect that, when we consider the important cases, the only interesting interpretation of "making an assumption" is such that it does entail having the corresponding belief. At least, when I was claiming that Barke's epistemic contextualism is incompatible with externalism I took "assumption" as entailing belief.

An important consequence of Barke's reliance on internalism is that epistemically circular arguments will never be available. Epistemically circular arguments have been described very well by William Alston, for example, in his book *The Reliability of Sense Perception* (1993). An epistemically circular argument is one where the *justification* (not the truth) of at least one of the premises depends on the *truth* of the conclusion. Such arguments are likely to come up in attempts to justify basic epistemic sources or methods. For example, we might try to argue for the reliability of sense perception by pointing out the favorable track record of sense perception. But our premise that sense perception has indeed a favorable track record can ultimately only be justified by using sense perception itself. So the justification of this premise presupposes the truth of the conclusion, the conclusion being the claim that sense perception is reliable. Similarly, one might try to argue for the reliability of induction by pointing to its favorable track record.¹ Again, the argument would turn out to be epistemically circular. Now, such epistemically circular arguments can never be used by an epistemic internalist, since under internalist assumptions epistemically circular arguments turn into logically circular arguments, i.e., arguments where the truth of at least one of the premises depends on the truth of the conclusion. This is so since, for an internalist, belief in the output of a certain method or source requires the subject to have the belief that the method or

source is reliable. And this is exactly what Barke has claimed too. But if the subject has to believe that the method or source is reliable in order to be justified in some output that is delivered by the method or source, then she can of course not use this output in order to justify the belief that the method or source is reliable. Thus, in effect, for the epistemic internalist epistemic circularity turns into logical circularity.² In contrast, an externalist does have the option of looking for epistemically circular arguments. They can provide genuine justification, as Alston has made very clear. For the internalist, such a kind of justification is excluded from the very beginning.³

This is especially worrisome since Barke is talking about a certain “intuition of circularity” and is trying to account for it. The intuition that she has in mind is the intuition that it is somehow circular to deduce *q* from *p*, if we have to assume that *q* in order to be justified or warranted in the belief that *p*. Now, this is indeed *logically* circular – *if* it is true that one has to assume that *q* in order to be justified or warranted in the belief that *p*. Barke and the epistemic internalist claim that this condition holds (in certain cases). But if one can be justified/warranted in believing that *p* without having to believe that *q* (and without having to be justified/warranted in believing that *q*), as the epistemic externalist holds (for these cases), then there is no logical circularity (and no other vicious circularity either). It may only be the case that *q* has to be *true* in order for one to be justified or warranted in one’s belief that *p*. And this only amounts to *epistemic* circularity – which is in itself not a vicious kind of circularity. The lesson then is: We have to be careful about what kind of circularity is involved, and epistemic circularity opens up the possibility of justification or warrant for the epistemic externalist in certain interesting cases where the internalist is doomed to vicious logical circularity.⁴

The second worry is that Barke’s version of contextualism is still very similar to conversational contextualism. I wonder whether epistemic contextualism deserves to be treated as a new kind of contextualism, and not simply as a revised version of conversational contextualism. In the following, I would like to explore the similarities and dissimilarities between epistemic contextualism and conversational contextualism.

There are two important similarities. The first one has to do with the topic of raising the standards. It is *contingent* whether someone makes a certain assumption, or challenges it. So, it is true that the questioning of an assumption is not, and need not be, epistemically motivated. (Similarly for conversational contextualism: it is contingent whether someone makes a “skeptical” statement that

raises the standards.) So the dynamics seem to be just as arbitrary, epistemically. From the epistemic point of view, the dynamics of context shifts are unexplained and unmotivated. Please note here that Barke never requires that the challenge of an assumption must itself be justified.

The second important similarity consists in the fact that there has to be something like the “mechanism” of forgetting a challenge for both positions. You cannot get around this, even as an epistemic contextualist, for the following reason. The assumption that sense perception is reliable has been challenged by philosophers, and is challenged frequently. But so far, it is fair to say, the challenge has not been successfully met. However, we are often allowed to make the assumption that sense perception is reliable. (Otherwise we could not use any empirical method, since any of them depends on the reliability of sense perception.) So there must be a way of silencing the challenge to sense perception (in order to avoid a far ranging skepticism). However, the only available way seems to be something like forgetting or ignoring the challenge. Providing an independent reason in favor of the assumption (which is the first legitimate way of answering the challenge) has proved terribly difficult. So if we are not willing to drop the assumption (the second legitimate way of reacting to the challenge), an option too hard to be swallowed, we simply have to do what the conversational contextualist recommends in this case: ignore or forget the challenge. There is no way around allowing such a mechanism, for both the conversational contextualist and the epistemic contextualist. Thus, it is not true to say, as Barke does, that the epistemic contextualist is better off in this respect. And both positions, therefore, face the problem of “upward stickiness of epistemic standards.” For how could we, especially as epistemologists, simply ignore or forget skeptical hypotheses again once they have ever bothered us?

Let us sum up the discussion of similarities. Both with respect to raising the standards and lowering them there is a fundamental similarity. The dynamics of context shifts are not, or need not be, genuinely epistemically motivated or explained. So Antonia Barke’s goal of providing a contextualist account with an epistemically motivated dynamics has not been reached. Let me add the following remark. Both kinds of contextualism exhibit a fundamental commonality, namely, they see knowledge as embedded within a *dialectical structure* of assumption and challenge, or default and challenge.⁵ The larger question that we have touched upon here, therefore, is whether knowledge is really to be conceived of as embedded in such a

default/assumption and challenge structure. The alternative is that knowledge is a non-dialectical phenomenon, a phenomenon that is simpler and more robust. Such an alternative account of knowledge has been given by Fred Dretske, Robert Nozick, and others. And in my view it is exactly the kind of problem that Barke has pointed out for the conversational contextualist that works equally against any dialectical conception of knowledge.

A third worry is concerned with a puzzling and counterintuitive consequence of Barke's epistemic contextualism (but also of conversational contextualism). Here, we have to consider the question of the overall plausibility of contextualism. Perhaps, what I will say here will be felt to be somehow question-begging. But it still seems to me to point to a crucial shortcoming of contextualism that should not be so easily dismissed. The point I wish to express is this: Making assumptions helps you to achieve knowledge, according to Barke's account. But how could *taking* certain conditions to be fulfilled be of any help in making the kind of *achievement* that knowledge consists in? How could this do any *real work*? The intuition behind this is of course that to gain knowledge is to make a genuine achievement, no matter what exactly this achievement consists in (whether knowledge requires sensitivity, safety, tracking, or something else like this). In contrast, it is intelligible how the *fact* that certain conditions (in the subject's environment) are fulfilled can help you to achieve something.⁶ But how could the mere making of the assumption that these conditions are fulfilled make any real contribution? What we need in order to arrive at knowledge are some positively supporting conditions, but the mere making of assumptions is no such support.

Perhaps another way of expressing the puzzlement would be putting it this way: the "naïve" person does make certain assumptions, and as a consequence of this, he knows; but the "more critical" person refrains from making these assumptions, and therefore she does not arrive at knowledge. This is a paradox of the type "paradox of epistemic laziness": doing less puts you in a better epistemic position. If contextualism has this implication, it seems to be off the track. Please note that this puzzlement is not special to Antonia Barke's version of contextualism but also attaches to conversational contextualism. Again, the intuition behind the worry is that the achievement of knowledge, this genuinely epistemic achievement, does not depend on those factors that the contextualists claim it to be dependent on. Whether someone succeeds in tracking the facts, or in acquiring a belief based on a conclusive reason, or in something like that, does not depend on whether someone makes a "skeptical" statement or brings

up some doubt about reliability or good working conditions. The evidence on which a certain belief is based has a certain epistemic strength, as we might say.⁷ It constitutes sensitivity, safety, tracking, or something like that. And whether it does so is independent of “skeptical” challenges. This, of course, presupposes that the parameters that enter into sensitivity, safety, and so on, are fixed objectively, independently of conversational moves and skeptical doubts in the mind of the knowing subject. But the comparison of the “naïve” and the “critical” person provides an intuition for exactly this claim.

The fact that we frequently do not know whether someone really has conclusive reasons for a belief, and that we change our mind about whether he does (depending, for example, on how important it is for us to find out the truth in question), does help to explain why we are often prepared to change our judgment about whether he knows, and why different persons differ in their judgment. So pointing towards such changes or differences in knowledge attributions is not an argument for contextualism.⁸ The phenomenon can be explained very well by the change or difference in our view about whether the conditions constitutive of knowledge are really fulfilled. Whether these conditions are fulfilled is independent of what we think, assume or know about it. Therefore, there is no reason for accepting the contextualist “paradoxical conclusion” that it might be true for one person to say “He knows that p” and at the same time true for another person to say “He does not know that p” (about one and the same subject). The difference in the knowledge attribution is due to a different view about whether the conditions constitutive of knowledge have been met. Since it is sometimes very difficult to tell whether the conditions have been met, it comes as no surprise that there is considerable instability in our disposition to attribute knowledge. But knowledge itself is a robust phenomenon.

ACKNOWLEDGEMENTS

Many thanks to Thomas Grundmann and Christoph Jäger for helpful discussions and comments.

NOTES

¹ For a more recent discussion of such an argument in favor of the reliability of induction see, for example, Lipton (2000).

² For further discussion of epistemic circularity see Hofmann (2001).

- ³ I have framed the discussion in terms of justification. If knowledge requires justification, everything applies equally well to knowledge. But even if warrant – what turns true belief into knowledge – is not a kind of justification, what has been said here is equally true of warrant.
- ⁴ According to Alston, we can go one step further: If we cannot decide by any independent means whether a certain basic source is reliable, we will go on believing that it is reliable (unless its outcomes undermine this belief). This, however, is practically rational and also provides a certain epistemic rationality, according to Alston. A criticism of this further view of Alston's can be found in Jäger (2002).
- ⁵ A default and challenge conception has been presented by Michael Williams, for example. Cf. Williams (2001). Please note that a dialectical structure in the sense intended here does not require linguistic utterances. For convincing criticisms of both Williams' position and contextualism in general, see Grundmann (2003).
- ⁶ This much has been made clear by externalists like Dretske. Cf. Dretske (1981).
- ⁷ Here I am drawing on ideas that have been put forward by Keith DeRose, of course without accepting DeRose's contextualist framework. Cf., for example, DeRose (1995).
- ⁸ Stewart Cohen has tried to argue for attributor contextualism along these lines in Cohen (2000). In his description of the case of Mary, John and Smith it is left open how reliable Smith's flight itinerary or the airline agent really is (Cohen, 2000, p. 92).

REFERENCES

- Alston, W.: 1993, *The Reliability of Sense Perception*, Cornell University Press, Ithaca.
- Barke, A.: 2004, 'Epistemic Contextualism', *Erkenntnis* **61**, 353–373.
- Barke, A.: 2002, *The Closure of Knowledge in Context*, Mentis, Paderborn.
- Cohen, S.: 2000, 'Contextualism and Skepticism', *Philosophical Issues* **10**, 94–120.
- DeRose, K.: 1995, 'Solving the Skeptical Problem', *Philosophical Review* **104**, 1–52.
- Dretske, F.: 1981, *Knowledge and the Flow of Information*, MIT Press, Cambridge, MA.
- Dretske, F.: 1970, 'Epistemic Operators', *The Journal of Philosophy* **67**, 1007–1023.
- Grundmann, Th.: 2003, 'Die Grenzen des erkenntnistheoretischen Kontextualismus', *Deutsche Zeitschrift für Philosophie* **51**, 1–22.
- Hofmann, F.: 2001, 'Können wir uns auf die Wahrnehmung verlassen?', in: Grundmann, Th. (ed.), *Erkenntnistheorie*, Mentis, Paderborn.
- Jäger, C.: 2002, 'Religious Experience and Epistemic Justification: Alston on the Reliability of "Mystical Perception"', in: C. U. Moulines and K.-G. Niebergall (eds.), *Argument und Analyse*, Mentis, Paderborn, 405–425.
- Lipton, P.: 2000, 'Tracking Track Records', *Proceedings of the Aristotelian Society* **74**, 179–205.
- Williams, M.: 2001, *Problems of Knowledge*, Oxford University Press, Oxford.

Philosophisches Seminar
 Universität Tübingen
 Bursagasse 1
 72070 Tübingen
 Germany
 E-mail: f.hofmann@uni-tuebingen.de

JOHN GRECO

A DIFFERENT SORT OF CONTEXTUALISM

ABSTRACT. A number of virtue epistemologists endorse the following thesis: Knowledge is true belief resulting from intellectual virtue, where *S*'s true belief results from intellectual virtue just in case *S* believes the truth *because S* is intellectually virtuous. This thesis commits one to a sort of contextualism about knowledge attributions. This is because, in general, sentences of the form “*X* occurred because *Y* occurred” require a contextualist treatment. This sort of contextualism is contrasted with more familiar versions. It is argued that the position: (a) yields a better solution to the lottery problem, and (b) may be grounded in a more general theory of virtue and credit.

1. INTRODUCTION

In very general terms, knowledge is non-accidentally true belief. A recent approach in epistemology tries to understand “non-accidentally” in terms of intellectual virtue.¹ The main idea is that, in cases of knowledge, *S* believes the truth not by accident, but because *S*'s belief is grounded in intellectual virtue. Hence, a number of virtue epistemologists have endorsed something like the following thesis:

Knowledge is true belief resulting from intellectual virtue.

Alternatively:

In cases of knowledge, *S* believes the truth *because S* believes out of intellectual virtue.

Here are some statements of the thesis in question.

Sosa: We have reached the view that knowledge is true belief out of intellectual virtue, belief that turns out right by reason of the virtue and not just by coincidence (Sosa, 1991, 277).

Zagzebski: It is important that success in reaching the end [i.e. believing the truth] is *due to* the other praiseworthy features of the act. The end must be reached *because of* those other features (Zagzebski, 1999, 107).

Riggs: [In cases of knowledge] the person derives epistemic credit ... that she would not be due had she only accidentally happened upon a true belief ... The



difference that makes a *value* difference here is the variation in the degree to which a person's abilities, powers, and skills are causally responsible for the outcome, believing truly that *p* (Riggs, 2002, 93–94).

Lehrer: [In cases of knowledge] the person is successful in accepting what is true *because* she accepts what she does in a trustworthy way in the particular case. Her trustworthiness explains her success in accepting what is true ... Her trustworthiness and the reliability of it explains her success in the particular case (Lehrer, 2000, 223).

Here we should understand 'virtue' in a broad sense: virtues are excellences of a sort. In the present context, the excellence is intellectual, and so a virtue is a cognitive ability or power. This follows Aristotle, who distinguishes intellectual virtues from moral virtues. The latter tend to be character traits, such as courage or temperance. The former tend to be powers, such as logical intuition and scientific demonstration. Of course, Aristotle thinks that prudence (a cognitive power) is both a moral and an intellectual virtue.

We may expand on Aristotle's list of intellectual virtues, so as to include intellectual excellences such as sound reason of various sorts, good memory, and accurate vision. The present thesis, then, amounts to this:

- (VK) In cases of knowledge, *S* believes the truth because her belief is produced by an intellectual excellence (a cognitive ability or power); for example, *S* reasoned well, or remembered correctly, or saw accurately.

What has gone largely unnoticed is that the present thesis commits one to a sort of contextualism about knowledge attributions. This is because, (a) according to VK, knowledge attributions involve causal explanations, and (b) in general, causal explanation language requires a contextualist treatment. Put another way, VK claims that knowledge attributions involve attributions of causal responsibility, and, in general, attributions of causal responsibility require a contextualist treatment.

Consider first the general point. In general, when we say that *Y* occurs "because" *X* occurs, or that *Y* occurs "by reason" of *X*'s occurring, we mark out *X*'s occurring as a particularly important or salient part of the causal story behind *Y*'s occurring. For example, to say that the fire occurred because of the explosion is not to say that the explosion caused the fire all by itself. Rather, it is to say that the explosion is a particularly important part, perhaps the most important part, of the whole story. Or to change the example: to say that the fire occurred because of *S*'s negligence is not to say that *S*'s negligence caused the fire all by itself. Rather, it is to say that *S*'s

negligence is a particularly salient part, perhaps the most salient part, of the set of relevant factors that caused the fire. But explanatory salience is a contextual matter: what is salient in the sense of “figuring importantly in an explanation” is partly a function of context.

What determines explanatory salience? Any number of things, but here we will consider two. First, among the various necessary parts of a complete causal process, an explanation will often pick out what is abnormal in the case. For example, we will say that sparks caused the fire if the presence of sparks in the area is not normal. That explanation misfires, however, if we are trying to explain the cause of a fire in a welding shop, where sparks are flying all the time. Or suppose that a white elephant walks into a room and causes a panic. Of course the white elephant entering the room is not sufficient all by itself to cause the panic – it would not if the room were part of a zoo and the people inside were animal trainers. But if the room is a place where white elephants are not normally found, and if the people inside *are* normal in their dispositions toward wild animals, we have no trouble picking out the elephant as “the” cause of the commotion.

Normality in this sense is an aspect of “subject context.” That is, it is an aspect of the context where the event in question (the subject of the causal explanation) is taking place. Other factors determining salience are a function of “conversational context,” or the context of the conversation in which a causal explanation is being offered, or causal responsibility attributed. An important aspect of this sort is our interests and purposes. For example, often when we are citing a cause we are citing something that we can manipulate to good effect. If the thing to be explained is smoke coming from the engine, for example, we will pick out some part that needs to be replaced. Here it is perfectly appropriate to say that the cause of the smoke is the malfunctioning carburetor, although clearly a faulty carburetor cannot cause smoke all by itself.

Another example: Sports fans will argue endlessly over why we lost the big game. Was it because we gave up too many points or because we didn’t score enough? Obviously, the outcome of a game is a function of both points allowed and points scored. The real argument here is over what was the most important factor in the loss. And *that* is a function of what one can reasonably expect, what could have been done differently, etc. For example, if we lose the soccer game by a score of ten–nine, we don’t blame the offense.

We may now turn to the more specific thesis that knowledge attributions involve causal explanations, or attributions of causal responsibility. According to VK, in cases of knowledge, *S* believes the

truth *because* *S* believes out of a cognitive ability or power. Or to use Sosa's formulation, *S* gets things right *by reason* of *S*'s cognitive abilities. But this is to say that *S*'s abilities are causally responsible for *S*'s getting things right. Alternatively, the fact that *S* believes out of intellectual virtue explains the fact that *S* believes the truth (as opposed to believing something false, or not believing anything at all).

Illustration

Let us call the resulting position 'virtue contextualism.' If virtue contextualism is correct, then there should be conversations involving knowledge attributions that manifest the contextualist elements here alleged. For example, there should be cases where disagreements over whether someone has knowledge reflect different emphases on the relative importance of *S*'s cognitive abilities (or virtues) to arriving at the truth. This is in fact the case. For example, consider the following conversation between a teenager and her mother.

Mother: How did you get home from the concert tonight, Honey?

Teenager: I got a ride home from a guy that I met there.

Mother: Are you crazy!? How did you know he wasn't an axe murderer?

Teenager: Don't be so dramatic, Mom. We talked for, like, an hour. I could see he was really nice.

Mother: I can't believe you! You're lucky you didn't end up on the side of the road somewhere.

Teenager: Mom! I think I'm a better judge of character than that.

The conversation can be interpreted as follows: The teenager is trying to take credit for her true beliefs, while the mother is trying to deny credit. Accordingly, the teenager makes *knowledge* claims: she judged correctly in this instance because, in general, she is a good judge about this sort of thing. The mother denies these knowledge claims, attributing her daughter's true beliefs to good luck rather than any good judgment on her part. Clearly, the mother and daughter have competing interests here, and this at least partly explains their disagreement regarding whether the daughter knows. In other contexts, however, interests and purposes might be more closely aligned.

Teenager: I met a great guy at the concert last night. He drove me home.

Friend: How did you know he was alright?

Teenager: We talked for, like, an hour. I could see he was really nice.

Friend: Cool.

Here the teenager's knowledge claim is readily accepted, perhaps because teenagers assume that they and their friends are excellent judges of character. In any case, there are no competing interests here, and so there is no pressure to attribute the teenager's true belief to anything other than her sound judgment.

Such conversations serve as a kind of anecdotal evidence that knowledge attributions are sensitive to context in the way suggested above. In the remainder of the paper I want to offer two more principled reasons in favor of this sort of contextualism. The first is that the present account of knowledge can be embedded in a more general theory of virtue and credit. In general, credit is grounded in virtue, and luck tends to undermine credit. In particular, intellectual credit is grounded in intellectual virtue, and luck tends to undermine intellectual credit. A second reason in favor of the present contextualist thesis is that it suggests an improved solution to the lottery problem.

2. A GENERAL THEORY OF VIRTUE AND CREDIT

The present approach claims that knowledge is true belief grounded in intellectual virtue. Put another way, in case of knowledge, *S*'s believing the truth can be put down to *S*'s intellectual abilities or powers. Now consider a common illocutionary force of knowledge attributions: Very often, when we attribute knowledge to someone we mean to give the person credit for getting things right. We imply that his believing the truth can be credited to him, as opposed to dumb luck, or blind chance, or something else. I want to suggest that these ideas are closely related. It is no accident that knowledge attributions often amount to a kind of credit attribution: In cases of knowledge, *S* deserves credit for believing the truth precisely because *S*'s believing the truth is grounded in *S*'s virtue, in *S*'s own abilities or powers. In the present section I want to explore these ideas and relate them to a more general theory of virtue and credit. We may begin by looking at some important work by Joel Feinberg concerning moral virtue and moral blame.

*Feinberg on Moral blaming*²

Feinberg's account of moral blaming takes off from the following central idea: When we attribute blame to a person for some occurrence, part of what we are doing is assigning causal responsibility to that person for the occurrence. Put another way, when we blame *S* for *X*'s occurring, we imply that *S* figures importantly into a correct causal explanation of why *X* occurred. For example, to blame someone for the fire is to imply that her actions caused it: she is the one who struck the match, or who did not pay attention, or who did pay the arsonist. Alternatively, we can blame a person for the action itself, implying that she herself was the action's cause, or perhaps that her choice was or her efforts were. As Feinberg notes, the distinction between blaming someone for her action and blaming someone for a consequence of her action is often merely verbal. For example, we can say either "She caused the fire by striking the match," or "She started the fire." Likewise, "She caused his death by poisoning his food," substitutes for "She killed him."

Second, Feinberg argues that when we blame someone for an action we imply that the action reveals something important about the person himself: "In general, I should think, a person's faulty act is registerable only if it reveals what sort of person he is in some respect about which others have a practical interest in being informed." (Feinberg 1970, p. 126). Feinberg's position is perhaps too strong on this point; it would seem that people can be rightfully blamed for actions that are out of character. Nevertheless, there does seem to be a *kind* of blame that Feinberg is right about. In other words, even if not all blaming implies that the person's action reveals a faulty character, there is a strong sort of blame, which is common enough, that does. Moreover, this strong sort of blame has a counterpart in a strong sort of credit. Often enough, credit for an action implies a judgment about the person as well, implying not only that the person is responsible for doing something good, but that this is a manifestation of virtuous character.

Putting all this together, Feinberg's account of blame for an action can be summed up as follows.

A person *S* is morally to blame for action *A* only if

- a.* *A* is a morally faulty action,
- b.* *A* can be ascribed to *S*, and

c. *A* reveals *S*'s faulty moral character.

Feinberg concludes that attributions of blame share the same pragmatics as causal explanations. His argument for this emphasizes clause (b) of the above account: attributing blame involves ascribing action, and ascribing action involves causal citation. What I want to emphasize, however, is that clause (c) acts the same way. Clause (c) also insures that attributions of blame involve causal citation, for what does it mean to say that an action reveals character, other than that the action results from character? In other words, clause (c) can be read:

(c') *S* did *A* because *S* has a faulty moral character.

This might seem too strong, and it is if we read (c') as saying that *S*'s character was sufficient all by itself to cause *S*'s action. Similarly, it is too strong if we read (c') as saying that, given *S*'s character, *S* had to do *A*. But it is not too strong if we remember the pragmatics of causal explanation language reviewed above. For according to that account, to say that *S*'s action is a result of her character is to say that *S*'s character is an important part, perhaps the most important part, of the story. Taken this way, (c') is not too strong at all, but rather reflects our common sense attitudes about the sources of human action. The fact is, we cite character in explanations of human behavior all the time, as when we say that he made the remark because he is insensitive (as opposed to having a bad day), or that she failed to spend the money because she is cheap (as opposed to hard up for cash at the moment).³ Feinberg's analysis reveals that such explanations are implied in attributions of blame, or at least in attributions of a certain sort of blame. And this implies that attributions of blame (of that special sort) will inherit the pragmatics of causal explanations. Clearly, any action will be the result of a number of factors, including a person's character. But sometimes we want to say that character is particularly salient – that it is an important part, perhaps the most important part, of the story behind why the person acted as he did.

A General Theory of Credit

Feinberg's account of moral blaming can easily be broadened in two ways. First, I have already noted that the counterpart of blame for an action is credit for an action. In fact, we can use credit as the general term, and talk about positive credit (i.e. praise) and negative credit

(i.e. blame) for an action. Second, there are kinds of credit other than moral.⁴ For example, we credit athletes for athletic feats and thinkers for intellectual ones. Accordingly, I propose the following as a general theory of credit.

A person *S* deserves credit of kind *K* for action *A* only if

- a.* *A* has value of kind *K*,
- b.* *A* can be ascribed to *S*, and
- c.* *A* reveals *S*'s *K*-relevant character. Alternatively: *S*'s *K*-relevant character is an important necessary part of the total set of causal factors that give rise to *S*'s doing *A*.

Two examples will illustrate this account.

Case A. Ken Griffey Jr. runs full speed toward the center field wall, leaps with outstretched glove, and catches the ball while diving to the ground. The home team crowd, just robbed of a game winning double, shakes their respective heads in admiration of Griffey's spectacular catch.

Case B. Griffey Jr. runs full speed toward the center field wall, trips, and falls face down on the ground. The ball bounces off his head, goes straight in the air, and comes down in his glove. The home team crowd, just robbed of a game winning double, shakes their respective heads in disgust.

In both cases, the action in question has clear athletic value – catching the ball before it hits the ground is essential to winning baseball games. Moreover, in both cases the catch is ascribable to Griffey – we can be sure that a broadcaster announcing the game will be yelling, “Griffey caught the ball! Griffey caught the ball!” But only in Case A will Griffey be given credit for catching the ball, and that is because in Case A Griffey's catching the ball is the result of his relevant character; i.e. his great athletic abilities. In Case B Griffey's catching the ball was just dumb luck, and so the home team crowd is not just a bunch of sore losers. They are right to be disgusted.

A similar phenomenon occurs when a poor fielder makes a spectacular catch. In this case he will be given credit of a sort – he will get pats on the back from his teammates and applause from the crowd. But it will be the same kind of credit that Griffey gets. *Griffey* makes spectacular catches all the time – *his* catches manifest his great skills. Not so when Albert Belle makes such a catch. If the catch is difficult, it is almost just good luck that he makes it. And opposing fans will treat it that way, withholding the credit they would readily give to Griffey.

A Theory of Intellectual Credit

When we attribute knowledge to someone we imply that it is to his credit that he got things right. It is not because the person is lucky that he believes the truth – it is because of his own cognitive abilities. He figured it out, or remembered it correctly, or perceived that it was so. Applying the account of credit attribution above, we have:

S deserves intellectual credit for believing the truth regarding *p* only if

- a.* believing the truth regarding *p* has intellectual value,
- b.* believing the truth regarding *p* can be ascribed to *S*, and
- c.* believing the truth regarding *p* reveals *S*'s cognitive abilities. Alternatively: *S*'s cognitive abilities are an important necessary part of the total set of causal factors that give rise to *S*'s believing the truth regarding *p*.

And hence: *S* knows *p* only if *S*'s believing the truth regarding *p* reveals *S*'s cognitive abilities. Alternatively: only if *S*'s cognitive abilities (or intellectual virtues) are an important necessary part of the total set of causal factors that give rise to *S*'s believing the truth regarding *p*.

Notice the equivalences (logical, not analytic):

S knows *p* iff *S* deserves credit for believing the truth regarding *p*.

S deserves credit for believing the truth regarding *p* iff *S* believes the truth regarding *p* because *S* is intellectually virtuous.

Hence:

S knows *p* iff *S* believes the truth regarding *p* because *S* is intellectually virtuous.

Comparison with Standards Contextualism

We may pause to compare the present account with standards contextualism. First, both accounts are contextualist; that is, they both make the truth conditions of knowledge claims relative to the context of attribution. But the way this works in the two accounts is different. According to standards contextualism, the context of attribution determines the standards for knowledge, so that standards are higher

or lower relative to different contexts. On the present account, however, the context of attribution determines the salience of various contributing causal factors, thus determining responsibility for true belief. Standards are not raised or lowered according to context; rather, responsibility for a complex event (someone's believing the truth) is creditable or not creditable to the believer according to context.

3. THE LOTTERY PROBLEM

Let us next consider how the present account helps to solve the lottery problem.⁵ The problem may be stated as follows. On the one hand, we want to say that there can be knowledge by inductive reasoning. On the other hand, it seems that a ticket holder does not know that she will lose the lottery, even if the odds are heavily in favor of her losing. So here is the problem: how is it that in general one can know through inductive grounds, but in the lottery case one fails to know, though one's inductive grounds are excellent?

To sharpen the problem, consider two cases of inductive reasoning.

Case 1. On the way to the elevator *S* drops a trash bag down the garbage chute of her apartment building. A few minutes later, reasoning on the basis of past experience and relevant background knowledge, *S* forms the true belief that the bag is in the basement garbage room. Of course her grounds for so believing are merely inductive: it is possible that the trash bag somehow gets hung up in the chute, although this is extremely unlikely.⁶

Case 2. *S* buys a ticket for a lottery in which the chances of winning are ten million to one. A few minutes later, reasoning on the basis of past experience and relevant background knowledge, *S* forms the true belief that she will lose the lottery. Of course her grounds for so believing are merely inductive: it is possible that she buys the winning ticket, although this is extremely unlikely.

Here is a third case, due to Jonathan Vogel:

Case 3. "Suppose two policemen confront a mugger, who is standing some distance away with a drawn gun. One of the officers, a rookie, attempts to disarm the mugger by shooting a bullet down the barrel of the mugger's gun. (I assume that the chances of doing this are virtually nil.) Imagine that the rookie's veteran partner knows what the rookie is trying to do. The veteran sees him fire, but is screened from seeing the result. Aware that his partner is trying something that is all but impossible, the veteran thinks (correctly as it turns out) [that the] rookie missed."

Many will have the intuition that *S* knows in Cases 1 and 3 but not in Case 2. But how so, given that her reasons are excellent in all the cases? This is what needs to be explained.

Contextualism and the lottery

Consider Cohen's own solution to the lottery problem. According to Cohen, the problem is solved by recognizing that attributions of knowledge are sensitive to context, and, more specifically, that the standards for knowledge are sensitive to context. We have knowledge in cases of ordinary inductive reasoning, such as that employed in the garbage chute case, because the standards that are operative in ordinary contexts are low enough to admit such cases as counting for knowledge. We do not have knowledge in the lottery case, however, because in that context the standards for knowledge are raised – the possibility of winning the lottery becomes salient, and our inductive evidence, as good as it is, does not rule out this possibility (Cohen, 1988, pp. 106–107).

I do not wish to deny Cohen's general point that the standards for knowledge are sensitive to context. It seems to me that they are. What is less clear is that standards contextualism helps to solve the lottery problem. Here are two reasons to think that it does not.

First, it is not clear why, on Cohen's account, *S* does not know that she will lose the lottery, even granting that the standards for knowing have been raised in the way that Cohen suggests. Cohen is quite explicit that he means to remain within the framework of fallibilism. Moreover, in the lottery case it is stipulated that *S* has excellent (although fallible) reasons for believing that she will lose. So why, on a fallibilist account of knowledge, does *S* fail to know that she will lose? To be clear, I am not claiming that *S* *does* know in the lottery case – I agree that she does not. My complaint is that nothing in Cohen's account explains *why* *S* does not know.

The same problem can be viewed from a different angle. Cohen says that when *S* reasons about the odds, the very form of her reasoning makes the possibility that *S* wins salient. And once made salient, Cohen says, that possibility cannot be ruled out. But again, why *can't it be*? Why is *S*'s reasoning about the odds good enough to rule out the possibility of winning, even once made salient? Again, it has been stipulated that *S* has *excellent* reasons for thinking she will not win the lottery, so why doesn't she know that she will not win? In sum, Cohen's contextualism does not explain what it was supposed to explain: given that we are fallibilists about knowledge, and given that we think inductive grounds are good enough to know in other cases, why are *S*'s grounds not good enough to know in the lottery case?

What are the prospects here for other versions of standards contextualism? The trick, of course, is for the standards contextualist to explain why *S* does not have knowledge in the lottery case, while at the same time preserving the intuition that *S* does have knowledge in other cases of inductive reasoning. But this will be hard to do. For example, Keith DeRose argues that *S* has knowledge if her belief matches the truth out to the nearest world where a salient alternative possibility is actual (see DeRose, 1995). However, the matching requirement insures that DeRose's account rules incorrectly in the garbage chute case and in the rookie cop case. This is because these cases are designed so that not-*p* worlds are very, very close. The world where the bullet goes down the mugger's gun barrel is nearly as close as worlds where it does not, although it is almost impossible that it does go down the barrel. Similarly in the garbage chute case: the weight and the trajectory of the bag has to be exactly right for the bag to get hung up, but the world where things are exactly right is very, very close, only minutely different from the way things are in the actual world. And so no matter how weak the standards for knowledge are being set, *S*'s belief in such cases will not match the truth far enough out into alternative possible worlds. The result will be that *S* does not know in the rookie cop case and the garbage chute case, no matter how weak the standards for knowledge in the context.

The second problem for standards contextualism is as follows. Suppose we grant that in the lottery case standards are raised so high as to undermine knowledge that *S* will lose the lottery. *S* utters the sentence "I don't know that I will lose," and her assertion is true relative to her own conversational context. Suppose also that, within the same conversational context, *S* claims that she knows where she parked her car. Is this assertion false? Or suppose that *S* says, "I know I bought a lottery ticket – here it is in my hand." Is that sentence false? Intuitively, *S* does know these things, even relative to her present context. But standards contextualism must say that she does not. According to standards contextualism, standards in *S*'s context have been raised so high as to make knowledge almost impossible. We got the result that *S* does not know that she will lose the lottery only because standards have been raised so high as to require infallibility, or perhaps near infallibility, for knowledge. But *S*'s beliefs about where she parked her car, or whether she bought a ticket, are not infallible either. And so it seems that *S* loses all her fallible knowledge in the present context, these items included. And that, I suggest, is highly implausible. Even if *S* does not know she will lose the lottery, surely she can know where she parked her car, or that

she bought a ticket! More exactly, surely her relevant knowledge claims are true, even relative to her conversational context.⁷

It is not open to the standards contextualist to deny that *S* knows she will lose the lottery, while at the same time insisting that she knows that she bought a ticket, or where she parked her car. True, we can imagine some mechanism of the conversational context that raises the standards for the one knowledge claim while keeping the standards low for the others. In this way, *S* would not be deprived of all her fallible knowledge in the context, even if she does not know she will lose the lottery. But that strategy is not open to standards contextualism. This is because: (a) it is an important part of that position that it preserves closure, and (b) closure is preserved only by holding that the same standards are applied to all knowledge claims made in the same conversational contexts. Standards contextualists such as Cohen and DeRose motivate their position precisely by showing how it preserves closure within conversational contexts. The standards contextualist is thus able to avoid “abominable conjunctions” such as “I do not know that I am a handless brain in a vat, but I do know that I have hands” (see DeRose, 1995). Closure is so achieved by holding that the standards for knowledge are determined by conversational context, in the sense that *all* knowledge claims in the same contexts get the same standards. It is this that insures that there is no single context relative to which *S* does not know the one thing but does know the other. Again, making standards depend on context, so that the same contexts get the same standards, is a non-negotiable part of the position.

The Lottery Problem Solved: Luck Undermines Credit

When we say that *S* knows *p*, we imply that it is not just an accident that *S* believes the truth with respect to *p*. On the contrary, we mean to say that *S* gets things right with respect to *p* because *S* has reasoned in an appropriate way, or perceived things accurately, or remembered things well, etc. We mean to say that getting it right can be put down to *S*’s own abilities, rather than to dumb luck, or blind chance, or something else. But then this gives us a resource for solving the lottery problem. For in the lottery case, it *does* seem to be just a matter of luck that *S* gets it right when *S* believes that she will lose the lottery. In the garbage chute case and the rookie cop case, however, we think that it is due to *S*’s good reasoning that she gets

things right – we give *S* credit for arriving at the truth in this case, and we are therefore willing to say that she knows.

The application of virtue contextualism to the lottery problem is therefore straightforward: knowledge attributions imply attributions of intellectual credit for true belief, and intellectual credit implies that the true belief is the result of *S*'s own intellectual abilities. But here as in other cases, salient luck undermines credit. In the lottery case, but not in the garbage chute case, it seems just a matter of luck that *S* believes the truth. In the garbage chute case, but not in the lottery case, *S*'s true belief is appropriately credited to her, i.e. to her intellectual abilities.

Of course, *S* employs admirable inductive reasoning no less in the lottery case than in the garbage chute case and the rookie cop case. In all these cases, therefore, *S*'s abilities make up a necessary part, but only a part, of the whole story regarding *S*'s believing the truth. But it is only in the garbage chute case and the rookie cop case that *S*'s abilities are a *salient* part of the story. In the lottery case, what is most salient is the element of luck.

Why does the element of luck become salient in the lottery case? I would suggest that the very idea of a lottery has the idea of chance built right into it, and chance is a form of luck. Here is the way we think of the lottery case: First, *S* reasons on the basis of excellent grounds that she will lose the lottery. Second, the lottery is held and reality either does or does not match up with *S*'s belief – it is just a matter of chance. Notice that things are different if *S* believes that she lost the lottery because she reads the results in the newspaper.⁸ Here again her evidence is merely inductive, but now the role of chance doesn't play an important part in the story. Here is the most natural way to think of the newspaper case: First the lottery is held and the facts are fixed. Second, *S* infers from a reliable source that she has lost the lottery. Now it is not just a matter of chance that she believes the truth – she believes the truth because she has the good sense to believe what she reads about the results.

To sum up the present section: I began by arguing that standards contextualism does not explain why *S* does not know that she will lose the lottery. First, the standards contextualist says that in the lottery case the standards for knowing get raised because the possibility of winning becomes salient. But this does not tell us why *S*'s reasoning fails to meet those standards, even if raised. That is, on the assumption that we want to remain fallibilists about knowledge, why doesn't *S*'s excellent fallible evidence allow her to know that she will lose the lottery? Second, standards contextualism implies that virtu-

ally all knowledge attributions are false relative to lottery contexts, since standards are raised so as to require infallibility (or close to it) for knowledge. But it seems wrong that one loses all fallible knowledge (so to speak) when thinking about the lottery. Perhaps we should see the two objections as constituting a dilemma for standards contextualism: Either fallible evidence is enough to know relative to lottery contexts or it is not. If it is, we have no explanation for why *S* does not know she will lose the lottery. If it is not, then we get the implausible result that, relative to lottery contexts, *S* does not know that she has bought a lottery ticket.

Virtue contextualism does much better in this respect. First, the very idea of a lottery involves the idea of chance, and so we have an explanation why chance (or luck) is salient in cases where the lottery is salient. We can then apply a familiar general principle of credit attribution to explain why *S* does not know that she will lose the lottery: namely, that salient luck undermines credit. Granted, there is much more to say here. But even before we say more, we have an independently motivated general principle that explains why *S* lacks knowledge that she will lose the lottery. Second, the explanation on the table does nothing to suggest that *S* lacks all fallible knowledge relative to the lottery context. For there is (typically) so salient luck corresponding to claims about where one's car is parked, or whether one has bought a lottery ticket. That is, the mechanism that robs *S* of knowledge that she will lose the lottery fails to rob *S* of other knowledge, even relative to the same context.

4. WHAT SORT OF LUCK UNDERMINES KNOWLEDGE?

The present solution to the lottery problem raises a pressing question: What sort of luck undermines credit, and hence knowledge? The problem is pressing because we do not want to say that *all* luck undermines knowledge. For example, it is only a matter of luck that I won the lottery, and that I am now enjoying an ocean cruise to celebrate. But that does not undermine my knowledge that dolphins are swimming along side the ship, as they do so in plain view.⁹ This is a problem even if we add that luck must be salient to undermine knowledge. For it is possible that I consider, just as I see the dolphins, how lucky I am to be here to see them.

Here is another example that needs explanation.¹⁰ Suppose that a soldier is charged with watching for enemy tanks crossing a valley. Almost always, the tanks are very well camouflaged and hence

difficult to spot. Moreover, our lookout is incompetent, almost never looking carefully enough to actually spot a tank. On this occasion, however, the driver of the tank is equally incompetent, and has forgotten to camouflage his vehicle. The lookout sees it in plain view and alerts his comrades that a tank is coming. Isn't it just good luck that the lookout sees the tank? Moreover, given the interests and purposes in place, the luck is salient. Yet we do not want to say that the lookout does not know that a tank is coming. Again, we need to say more about what sort of luck undermines knowledge.

A suggestion by Mylan Engel gives us help here. Engel makes a distinction between "evidential luck" and "verific luck." Roughly, evidential luck is luck regarding what evidence one has. Verific luck is luck regarding whether one believes the truth, *given that one has the evidence that one does*.¹¹ Engel argues that verific luck undermines knowledge whereas evidential luck does not. And indeed, this suggestion handles the present cases well. It seems to me that Engel's distinction is not entirely adequate, however, since it might be that some kinds of knowledge do not need evidence. For example, it seems that some kinds of introspective knowledge and some kinds of memory knowledge do not. If this is right, then we need a more general way to mark the relevant distinction. Perhaps it is this: Luck regarding how one gets into one's present situation does not undermine knowledge. Call this sort of luck "situational luck." Verific luck can then be understood as luck regarding whether one believes the truth, *given that one is in the situation that one is*. As before, verific luck does undermine knowledge. This way of drawing the distinction continues to handle the present examples well, but allows the possibility that some knowledge does not involve evidence.

The present point generalizes to other kinds of credit. In general, luck in the way that one gets into one's situation does not undermine credit for what one does, once in that situation. This is one of the lessons of the literature on moral luck.¹² A general prohibition on luck shrinks the sphere of moral responsibility to nothing, and hence an adequate theory of moral responsibility must allow for the influence of some kinds of luck but not others. Plausibly, one can be morally responsible for the way one acts in a situation, even if one's being in that situation in the first place is almost entirely the result of good (or bad) luck. Similarly, one gets credit for making a great catch in the World Series, even if it is just good luck that one is in the World Series to begin with.

ACKNOWLEDGMENTS

I would like to thank Peter Baumann, Michael Bergmann, Dirk Koppelberg, and Mark Lance for their comments on earlier drafts and in conversation.

NOTES

- ¹ For example, see Sosa (1991), Zagzebski (1999), Lehrer (2000) and Greco (2003). The present paper draws on material from Greco (2003).
- ² Feinberg's discussion takes place over three papers, all of which are collected in Feinberg (1970). The papers are "Problematic Responsibility in Law and Morals," "Action and Responsibility," and "Causing Voluntary Actions." Page numbers that follow correspond to Feinberg (1970).
- ³ Some recent work in social psychology suggests that common sense is flawed in this respect. For example, see Ross and Nisbett (1991). For a persuasive argument against such a conclusion, see DePaul (forthcoming).
- ⁴ Feinberg's own discussion is at times aimed at other kinds of blame.
- ⁵ For a nice statement of the problem, see Cohen (1988), Cohen discusses the problem again in Cohen (forthcoming).
- ⁶ The example is from Sosa (2000), We can imagine that the bag's getting hung up is extremely unlikely because everything would have to go just right for that to occur, including the trajectory of the bag, its contents, the distribution of its weight, etc.
- ⁷ Similar considerations are raised by Peter Baumann in the paper he wrote for this volume.
- ⁸ This point is made by Cohen in Cohen (1988).
- ⁹ The example is fictitious.
- ¹⁰ The example was suggested to me by Mark Lance in conversation.
- ¹¹ See Engel (1992). I thank Michael Bergmann for pointing out to me that Engel's distinction is helpful in the present context.
- ¹² See Nagel (1979), Walker (1991), and Greco (1995).

REFERENCES

- Cohen, S.: 1988, 'How to Be a Fallibilist', *Philosophical Perspectives* **2**, 91–123.
- Cohen, S.: forthcoming, 'Contextualist Solutions to Epistemic Problems: Skepticism, Gettier and the Lottery', *Australasian Journal of Philosophy*.
- DePaul M. and Zagzebski L. (eds.): 2003, *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Oxford University Press, Oxford.
- DePaul, M.: forthcoming, 'Character Traits, Virtues and Vices: Are there None?', in *Proceedings of the Twentieth World Congress of Philosophy, Volume IX: Philosophy of Mind and Philosophy of Psychology*.
- DeRose, K.: 1995, 'Solving the Skeptical Problem', *Philosophical Review* **104**, 1–52.
- Engel, M. Jr.: 1992, 'Is Epistemic Luck Compatible with Knowledge?', *The Southern Journal of Philosophy* **XXX**, **2**, 59–75.

- Feinberg, J.: 1970, *Doing and Deserving: Essays in the Theory of Responsibility*, Princeton University Press, Princeton.
- Greco, J.: 1995, 'A Second Paradox Concerning Responsibility and Luck', *Metaphilosophy* **26**, 81–96.
- Greco, J.: 2003, 'Knowledge as Credit for True Belief', in DePaul M. et al. (eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Oxford University Press, Oxford.
- Lehrer, K.: 2000, *Theory of Knowledge*, 2nd edn. Westview Press, Boulder, CO.
- Nagel, T.: 1979, *Mortal Questions*, Cambridge University Press, Cambridge.
- Riggs, W.: 2002, 'Reliability and the Value of Knowledge', *Philosophy and Phenomenological Research* **64**, 79–96.
- Ross, L. and R. Nisbett (1991), *The Person and the Situation*, McGraw-Hill, New York.
- Sosa, E.: 1991, *Knowledge in Perspective*, Cambridge University Press, Cambridge.
- Sosa, E.: 2000, 'Skepticism and Contextualism', *Philosophical Issues* **10**, 1–18.
- Walker, M.: 1991, 'Moral Luck and the Virtues of Impure Agency', *Metaphilosophy* **22**, 14–27.
- Zagzebski, L.: 1999, 'What Is Knowledge?', in Greco J. and Sosa E. (eds.), *The Blackwell Guide to Epistemology*, Blackwell Publishers, Oxford, 92–116.

Department of Philosophy
 Fordham University
 Bronx, NY 10458
 USA
 E-mail: greco@fordham.edu

DIRK KOPPELBERG

ON THE PROSPECTS FOR VIRTUE CONTEXTUALISM:
COMMENTS ON GRECO

ABSTRACT. John Greco has proposed a new sort of contextualism which exhibits a principled grounding in an agent reliabilist virtue epistemology. In this paper I will discuss Greco's two main reasons in favor of virtue contextualism. The first reason is that his account of knowledge can be derived from a more general theory of virtue and credit. The second reason consists in the thesis that a virtue contextualist solution to the lottery problem is superior to standards contextualism. With regard to the first claim, I raise some questions concerning the status and the content of the crucial conditions for Greco's theory of intellectual credit. With regard to the second claim, I try to show that his arguments do not succeed in establishing the superiority of virtue contextualism to standards contextualism. I close with some remarks on the relation among Greco's virtue contextualism, the traditional approach to the theory of knowledge and the proper domain of contextualism.

1. INTELLECTUAL VIRTUE, VIRTUE EPISTEMOLOGY AND VIRTUE
CONTEXTUALISM

John Greco has proposed a new sort of contextualism which exhibits a principled grounding in an agent reliabilist virtue epistemology. The common core of a virtue epistemologist's approach to propositional knowledge consists in the thesis that knowledge is true belief resulting from intellectual virtue. But what is intellectual virtue? This is itself a hotly debated question among different proponents of current virtue epistemology. In very general terms, maximalist or responsibilist accounts on the one hand are distinguished from minimalist or reliabilist accounts of intellectual virtue on the other hand. According to a maximalist or responsibilist account, a state of a person is a virtue if and only if it is (i) acquired, (ii) stable, (iii) an excellence that allows its possessor to reliably succeed in realizing a certain end, and (iv) a motivation to realize an end suitable to the end in (iii). Reliabilists have doubted that all virtues, and especially the virtues which are supposed to be necessary for knowledge, have to be acquired and have to include a motivation to bring about the desired end. In their view, maximalists are misled by Aristotle's treatment of ethical virtues (*ethikai aretai*) when they insist that intellectual virtues (*diano-etikai aretai*) must be acquired traits which include a strong



motivational component. For agent reliabilists like Sosa and Greco the appropriate explication of intellectual virtues can be traced back to a broader Platonic sense of virtue, according to which anything with a function, such as a knife or an eye, has a virtue germane to it. Reliabilists understand intellectual virtues as truth-conducive stable cognitive dispositions, abilities or powers, including genetically-endowed ones such as the various perceptual faculties and the transmission faculties of memory and reasoning. Such a view is the foundation for Greco's point of departure which he formulates in the following thesis:

VK In cases of knowledge, *S* believes the truth because her belief is produced by an intellectual excellence (a cognitive ability or power); for example, *S* reasoned well, or remembered correctly, or saw accurately (Greco, 2004, p. 384).

According to Greco, this thesis VK has two further theses VC in its train which I regard as the most important claims of his view and which may be put as follows:

VC First, a virtue epistemologist approach to knowledge commits one to a sort of contextualism about knowledge attributions. Second, this sort of contextualism, which may be called "virtue contextualism," is superior to standards contextualism in solving important philosophical problems.

In the following I will focus on Greco's two main reasons in favor of virtue contextualism. The first reason maintains that his account of knowledge can be derived from a more general theory of virtue and credit. The second reason consists in the thesis that his account offers an improved solution to the lottery problem. With regard to the first claim, I raise some questions concerning the status and the content of the crucial conditions for the theory of intellectual credit. With regard to the second claim, I try to show that Greco's arguments fail to establish the superiority of virtue contextualism to standards contextualism. I will close with some remarks on the relation among Greco's virtue contextualism, the traditional approach to the theory of knowledge and the proper domain of contextualism.

2. VIRTUE CONTEXTUALISM AND INTELLECTUAL CREDIT

Greco derives his virtue contextualist account of knowledge from a more general theory of virtue and credit which in turn is a reconstructive generalization of Joel Feinberg's account of moral blaming

(Feinberg 1970). Feinberg's account is based on two central ideas. The first idea deals with an adequate explication of blaming. When we blame a subject *S* for some occurrence *X*, we presuppose that *S* figures importantly into a correct causal explanation of why *X* occurred. The correct causal explanation singles out one of the relevant causal conditions that is especially interesting to us, given our special cognitive concerns and our various practical purposes. A judgment that cites one of the numerous eligible causal conditions for an event as "the cause" Feinberg calls a causal citation. A causal citation in terms of salient features is partly a function of context. Feinberg's second idea maintains that there is a strong sort of blame that reveals something important about the subject's faulty moral character. In this case we do not blame a person just for a certain action or for a consequence of her action. We blame that person because her action illuminates something important about the person herself: her action reveals her character. Putting these two ideas together, we can conclude that for a certain sort of blame attribution character is particularly salient, that it is, as Greco puts it "an important part, perhaps the most important part, of the story behind why the person acted as he did" (Greco, 2004, p. 389).

On the basis of these central ideas, Greco summarizes Feinberg's account of blame for an action as follows (Greco, 2004, p. 388):

A person is morally to blame for action *A* only if

- a. *A* is a morally faulty action,
- b. *A* can be ascribed to *S*, and
- c. *A* reveals *S*'s faulty moral character.

For Greco to say that an action reveals character means the same as that the action results from character and hence he reads clause (c) as

c'. *S* did *A* because *S* has a faulty moral character.

Of course, *S*'s character is not sufficient by itself to cause *S*'s action instead of which it is crucial that *S*'s character is an important necessary part of the total set of causal factors in the story.

Taking into account that the counterpart of blame for an action is credit for an action and that there is not only a moral kind of credit, Greco proposes the following as a general theory of credit (Greco, 2004, p. 390):

A person *S* deserves credit of kind *K* for action *A* only if

- a. *A* has value of kind *K*,

- b. *A* can be ascribed to *S*, and
- c. *A* reveals *S*'s *K*-relevant character. Alternatively: *S*'s *K*-relevant character is an important necessary part of the total set of causal factors that give rise to *S*'s doing *A*.

From this general theory of credit Greco derives his theory of intellectual credit in the following way:

S deserves intellectual credit for believing the truth regarding *p* only if

- a. believing the truth regarding *p* has intellectual value,
- b. believing the truth regarding *p* can be ascribed to *S*, and
- c. believing the truth regarding *p* reveals *S*'s cognitive abilities. Alternatively: *S*'s cognitive abilities are an important necessary part of the total set of causal factors that give rise to *S*'s believing the truth regarding *p*.

And hence:

S knows *p* only if *S*'s believing the truth regarding *p* reveals *S*'s cognitive abilities. Alternatively: only if *S*'s cognitive abilities (or intellectual virtues) are an important necessary part of the total set of causal factors that give rise to *S*'s believing the truth regarding *p*.

In the following I want to discuss three items. The first of them concerns the logical status of the proposed definitions. The second deals with the supposed analogy of condition (b) in Feinberg's and Greco's accounts. The third focuses on the connection between cognitive abilities and cognitive character in condition (c).

First, and taken literally, the introduced accounts of blaming and credit formulate only individually necessary but not jointly sufficient conditions for their respective subjects. I wonder whether this is just a slip or whether there is a deeper reason behind it. If the second possibility should turn out to be true, it would be good to have Greco's comment on this issue, because as things stand right now, we just do not have a complete account of intellectual credit which can do the job that it is supposed to do within a virtue theory of knowledge. It goes without saying that we need a complete set of conditions for the concept of deserved credit if this concept is regarded as crucial for a successful explication of the concept of knowledge.

With regard to the second point, condition (b) formulating an ascription clause plays an important role within Feinberg's account of moral blaming. A person *S* is morally to blame for action *A* only if *A* can be ascribed to *S*, so that attributing blame involves ascribing action, and ascribing action involves causal citation in terms of salient features. Turning to Greco's account of intellectual credit,

one may wonder what the ascription clause is meant to involve in the domain of belief. Believing the truth regarding p can hardly be understood as an action so that the analogy has its limits. Concerning the ascription of believing the truth to S , Jonathan Kvanvig has recently complained that “it is hard to see what it could involve beyond the simple claim that the person believes the claim in question. That is, under what conditions might we say that a person has a true belief but that the true belief cannot be ascribed to that person? The claim that there is some distinction here escapes me, and hence I see no reason whatsoever for distinguishing the two.” (Kvanvig, 2003, p. 89) Nevertheless there is an important reason for distinguishing the two, and it has to do with the demand that the ascription clause involves the causal citation in terms of salient features which are the subject’s own cognitive abilities. Under the condition that the subject is just lucky to believe the truth, the true belief would not be *ascribed* to that person. So to avoid any further misunderstanding, condition (b) within a theory of intellectual credit should better be formulated with explicit reference to S ’s own cognitive abilities.

But now we arrive at a third problem, because it is no longer obvious what distinguishes condition (b) from condition (c) which formulates that believing the truth regarding a certain proposition reveals S ’s cognitive abilities. Both in Feinberg’s account of moral blaming and in Greco’s reconstruction of a general theory of credit the focus is not just on abilities, but on character. This is significant because the grounding of abilities or powers in character gives virtue epistemology its special outlook. One may have a power or ability one infrequently exercises if at all. A virtue’s nature is much more like that of a habit. It is a matter of what one would do under certain conditions rather than of what one is able to do. And so the relevant cognitive abilities have to be embedded in S ’s reliable cognitive character to elucidate what is important about intellectual virtues.¹ When the difference between cognitive abilities and their integration into a cognitive character is neglected, the crucial distinction between mere cognitive abilities and intellectual virtues is blurred. The basic feature of virtue epistemology gets lost.

A question that comes up immediately after this clarification is the following: Why should it in general be necessary for knowledge that S ’s believing the truth regarding p has to be grounded in S ’s cognitive abilities understood as his stable character dispositions? One might ask why the disposition that results in true belief has to be part of S ’s cognitive character and why in addition to being a reliable

disposition, the disposition has to be stable. Imagine that quite often Henry is a sloppy and superficial reasoner, so that he seems to lack the intellectual virtue of careful and circumspect reasoning. But there are other occasions on which he reasons thoroughly and validly. Should we deny that on these occasions his blameless intellectual behavior can give him knowledge? Is it not conceivable or even plausible to have a power one infrequently exercises if at all? Or, to put these questions slightly different and in more general terms: Is the stable ability or disposition to carry out belief forming processes necessary for an acceptable account of knowledge?

If we are inclined to give a negative answer to this question, we understand the relevant sense in which dispositions are stable, as stability through time. However, in a recent reply to his critics, Greco has made it unmistakably clear that the proper sense in which dispositions are stable, must be understood in terms of what a subject does in relevantly close possible worlds (cf. Greco, 2003b, p. 472 f.). The basic idea is that the further out into logical space the relevant behavior persists, the more stable the disposition. So if it is not just an accident that Henry reasons on some occasions thoroughly and validly, but produces a high rate of success in those conditions over the range of relevantly close worlds, we will say that he knows, even if he does not have the ability to reason rigorously all the time he is thinking things through.

But now one might complain that even this kind of defense is not sufficient for solving the much discussed “problem of strange and fleeting processes” such as BonJour’s famous clairvoyant who forms true beliefs based on clairvoyance even though he knows better or Plantinga’s serendipitous brain lesion that happens to cause a true belief that such a lesion is present. We have seen that according to Feinberg, ascribing moral credit for a subject’s action has to be grounded in his moral character. According to Greco, ascribing intellectual credit for a subject’s belief has to be grounded in his cognitive character (even if he no longer puts it this way in his definition). But a cognitive character is not just a collection of stable cognitive dispositions, at least equally important is a high degree of cognitive integration of the relevant abilities. I think Greco is on the right track when he claims that it is due to the lack of cognitive integration of both the processes operative in BonJour’s and in Plantinga’s cases that we refrain from ascribing knowledge to the respective subjects in the examples.

More than thirty years ago, Karl Popper presented a famous paper “Epistemology Without a Knowing Subject.” If we are attracted

to an agent reliabilist virtue theory, we will pursue an epistemology *with* a knowing subject in a very strong sense. Our idea and definition of knowledge will be conceptually dependent on a detailed account of what it means to have an integrated cognitive character. What virtue epistemologists are looking for and what they have to develop, may turn out as something like a coherence theory of cognitive abilities. Unlike a coherence theory of beliefs, the relevant relations among cognitive abilities cannot be purely inferential. So is there any philosophically satisfying way to characterize the relations Greco is looking for? A second point concerns my guess that any psychologically plausible cognitive integration will be a matter of degree so that the integration of cognitive abilities needed for knowledge will become a gradual matter, too. But then knowledge itself seems to become a gradual good in a very sophisticated sense: in each case of knowledge attribution we have to establish the stability of the relevant cognitive abilities and the measure in which they cooperate and interact with other aspects of the cognitive system. Hence, virtue epistemologists face an important shift of inquiry: in analyzing knowledge they end up with the challenging task to give an informative and detailed account of intellectual character.

3. VIRTUE CONTEXTUALISM AND THE LOTTERY PROBLEM

Greco defends his virtue contextualist account of knowledge by claiming that it solves the lottery problem. More precisely, he claims that virtue contextualism is superior to standards contextualism in solving this problem. The problem consists in the following question: how is it possible that generally one can gain knowledge through inductive grounds whereas in a fair lottery a ticket holder does not know that she will lose, though her inductive grounds for losing are excellent?

To bring the problem into focus, let us recapitulate two of Greco's three cases of inductive reasoning (Greco, 2004, p. 392):

Case 1. On the way to the elevator *S* drops a trash bag down the garbage chute of her apartment building. A few minutes later, reasoning on the basis of past experience and relevant background knowledge, *S* forms the true belief that the bag is in the basement garbage room. Of course her grounds for so believing are merely inductive: it is possible that the trash bag somehow gets hung up in the chute, although this is extremely unlikely.

Case 2. *S* buys a ticket for a lottery in which the chances of winning are ten million to one. A few minutes later, reasoning on the basis of past experience and relevant background knowledge, *S* forms the true belief that she will lose the lottery. Of course her grounds for so believing are merely inductive: it is possible that she buys the winning ticket, although this is extremely unlikely.

Why are so many people inclined to admit that subject *S* knows that her trash bag falls down the garbage chute and reaches the basement room whereas they deny that *S* knows that she will lose the lottery? The challenge consists in explaining why *S* does *not* know in the lottery case, while at the same time preserving the intuition that *S* *does* know in the garbage chute case.

For the solution of this problem a standards contextualist like Stewart Cohen appeals to the insight that standards for knowledge are sensitive to context (cf. Cohen, 1988, pp. 106–108). In the garbage chute case the standards of ordinary inductive reasoning are low enough to regard such cases as knowledge. In the lottery case, however, the standards for knowledge are raised because the possibility of winning the lottery becomes salient and even our excellent inductive evidence cannot rule out this possibility.

Even if Greco agrees with Cohen's general point that the standards for knowledge are context-sensitive, he doubts that standards contextualism is able to solve the lottery problem. In the following I will discuss his two reasons for thinking that it will not.

The first point concerns his suspicion that nothing in Cohen's account explains why *S* does not know that she will lose the lottery. Is Greco's suspicion warranted?

If I understand Cohen correctly, he would stress that the reason consisting in the $R - 1/n$ probability that the ticket will lose does not entail that *S* in fact loses. Obviously, there exists the alternative that the ticket will win. Since we would deny that *S* knows that she loses, Cohen concludes that this alternative is relevant in this context.

Of course, the crucial question is what makes it relevant. According to Cohen, the explanation lies in the statistical nature of the reasons:

Although, as fallibilists, we allow that *S* can know *q*, even though there is a chance of error [...], when the chance of error is salient, we are reluctant to attribute knowledge. Statistical reasons of the sort that *S* possesses in the lottery case make the chance of error salient. The specification that *S*'s reasons is the $n-1/n$ probability that the ticket loses, calls attention to the $1/n$ probability that the ticket wins. Our attention is focused on the alternative that the ticket wins and this creates a context in which we are reluctant to attribute knowledge, unless *S* has some independent ground sufficient for denying the alternative. (Cohen, 1988, p. 106)

Cohen's point is that the very form of statistical reasoning makes the possibility salient that *S* wins and once made salient, this possibility cannot be eliminated. Greco does not regard this point as an acceptable explanation because he wonders why *S*'s reasoning about the odds is not good enough to rule out the possibility of winning, even once made salient. I think, Cohen might answer, that once made salient, standards arise to a point where *S*'s original reasoning about the odds just does not give us the focus for what we are looking for right now.

Cohen's solution has to do with his claim that after our attention has been drawn to the $n-1/n$ probability that the ticket will lose, this in turn calls our attention to the $1/n$ probability that the ticket will win and hence our attention is focused on it as a relevant alternative. Interpreted in a descriptive way concerning what will actually happen when people are confronted with the prospects of the lottery, I think this claim is at least dubious. Now I suppose that Cohen does not want to understand it descriptively. Interpreted normatively, one might say, that when there is a $1/n$ probability that the ticket will win, this ought to be salient for the reasoner in this particular situation. And when it is salient and therefore turns out to be a relevant alternative, our initial inductive reasoning is no longer sufficient to rule this alternative out so that we would be entitled to claim that *S* knows that she will lose.

Greco has put forward a different solution to the lottery problem. The basic idea is that in contrast to the garbage chute case, in the lottery case it seems to be just a matter of chance that *S* gets it right when she forms the belief that she will lose the lottery. So in the lottery case the element of chance becomes salient. Greco's explanation why chance is salient in cases where the lottery is salient, consists in his suggestion that the very idea of a lottery conceptually involves the idea of chance. But in exactly which sense does this provide us with a better explanation than the one given by standards contextualism? Until we get an answer to this question, I do not think that the second part of Greco's initial thesis VC has been proven.

Consider what a standards contextualist might say in comparing the lottery case and the garbage chute case. The very idea of a lottery involves the relevant possibility to win and this relevant possibility has to be excluded for knowing that the ticket will lose. The garbage chute case is told in a way that it does not involve the relevant possibility that the trash bag somehow might get hung up in the chute and so people will say that they know that the bag is in the basement garbage room.

For Greco, the element of luck is salient. For Cohen, the chance of error is salient. Virtue contextualists appeal to the general principle that chance undermines credit; standards contextualists might appeal to the general principle that error undermines knowledge. Both accounts have interesting resources for drawing the important distinction between the lottery case and the garbage chute case. Hence, we need additional reasons for claiming that one of them is superior to the other.

I think that Greco tries to deliver such a reason with his second problem for standards contextualism.

Suppose we grant that in the lottery case standards are raised so high as to undermine knowledge that *S* will lose the lottery. *S* utters the sentence “I don’t know that I will lose,” and her assertion is true relative to her own conversational context. Suppose also that, within the same conversational context, *S* claims that she knows where she parked her car. Is this assertion false? [...] Intuitively, *S* does know these things, even relative to her present context. But standards contextualism must say that she does not. (Greco, 2004, p. 394)

Unfortunately, I see no reason why *S* *does* actually know where she parked her car *even* relative to the *present* context. Because we tend to think of the case in terms of quite restricted scenarios, we are not inclined to consider certain alternatives as relevant and thereby change the initial standards. But they may change if *S* is reminded of the fact that *S* had quite often claimed to know where she parked her car but nonetheless was wrong. So when Greco maintains that “surely she can know where she parked her car” (Greco, 2004, p. 394) I’m afraid that this example is just begging the question against standards contextualism. When it is the same conversational context, in which we deal with the lottery case and the parked car case, standards are raised so high that both the knowledge that *S* will lose the lottery and the knowledge where *S* parked her car will be undermined.

At this point it might be objected that the problem with Greco’s argument is not that he begs the question, but that he might be overlooking ways in which standards contextualists can avoid taking the position he attributes to them.² According to such an objection, contextualists have the option of saying that standards are raised with regard to knowing that *S*’s ticket will lose in virtue of the statistical nature of error whereas standards are not raised with regard to the nonstatistical question of whether *S*’s car is still where she parked it earlier. So one may doubt that a contextualist like Cohen would agree with Greco that, since the standards are raised with regard to lottery tickets, they are also raised with regard to parked cars. But if stan-

dards are not raised in both cases, Cohen could say that whereas *S* does not know that her ticket is a loser, she does know where she parked her car.

Unfortunately, this proposal turns out to be a non-starter because the strategy is not open to standards contextualism, as Greco rightly emphasises (cf. Greco 2004, p. 394). It is a defining mark of that position that it wants to preserve closure. And if closure is preserved only by applying the same standards to all knowledge claims made in the same conversational contexts, a standards contextualist just does not have the option to raise standards just in one case and not in the other. So I do not see a chance for such a proposal to succeed in defending standards contextualism against Greco's charge.

To sum up this section: Greco claims that the explanation of the lottery problem lies in the statistical nature of the lottery. Cohen claims that the explanation lies in the statistical nature of our reasons. Is there a definitive explanatory advantage on either of the two sides? What needs further clarification is the general relation between standards contextualism and virtue contextualism. For standards contextualists, the truth-value of a knowledge ascription is sensitive to certain facts about the speaker and the hearer of the context. For virtue contextualists, the relation between the necessary conditions of true belief and credit is spelled out via a sort of causal explanation that involves contextual parameters. These seem to be rather different contextualist accounts, and at least for me, it is not quite clear in which way they are related. It is even less clear whether we are warranted in claiming that one of them supersedes the other with regard to the solution of the lottery problem.

4. VIRTUE CONTEXTUALISM IN CONTEXT

I propose the following account as a reconstruction of Greco's argument for virtue contextualism.

- AVC*
1. Knowledge attributions require attributions of intellectual credit.
 2. Attributions of intellectual credit require a sort of causal explanation.
 3. The appropriate causal explanation requires causal citation in terms of salient features.
 4. Causal citation in terms of salient features is partly a function of context.
 5. Therefore, knowledge attributions require a partly contextualist causal citation explanation.

Or, to put the conclusion shortly (and more vaguely):

Knowledge attributions require a partly contextualist treatment.

In some of his papers Greco writes that a virtue epistemologist's definition of knowledge commits one to a sort of contextualism about knowledge or knowledge attribution. So does he give us an account of knowledge? Or is it supposed to be an account of the proper use of "knowledge"? Or does Greco regard these two options as a distinction without a difference? I do not think so, because in a recent paper we read the following:

We may distinguish two questions one might try to answer when giving an account of knowledge. The first is the "What is knowledge?" question. This question asks what conditions a person must satisfy to count as knowing. The second is the "What are we doing?" question. This question asks what illocutionary act is being performed when we say that someone knows. (Greco, 2003a, p. 116)

Now I wonder whether it is advisable to unite the answers to these two questions within a single definition as Greco finally proposes. Of course, one can argue that a complete epistemology must involve two components, one about the knowledge attributor and one about the property referred to. The attributor uses the term "knowledge"; the knowing subject possesses the property that is referred to in that use. Contextualism may turn out to be the right theory about the use of the term. A theory of knowledge, at least as it is traditionally conceived, is a theory about the epistemic property in any given context. Now virtue contextualism seems to strive for a theory both about the conditions a person must satisfy to count as knowing and about the illocutionary acts being performed when ascribing knowledge to someone.

When we think about a contextualist treatment of knowledge or knowledge attribution, we have to consider the familiar distinction between attributee based contextualism and attributor based contextualism. What is their relation to a reliabilist virtue epistemology? It goes without saying that each form of reliabilism is committed to attributee based or modest contextualism because modest contextualism coincides with externalism. According to Greco, a reliabilist virtue epistemology is also at least consistent with an attributor based contextualism in form of standards contextualism: both the degree and the range of reliability required for knowledge ascriptions change with the attributor context. Virtue contextualism goes beyond standards contextualism in claiming that knowledge attributions require attributions of intellectual credit which in turn require a sort of causal explanation. The appropriate causal explanation requires causal

citation in terms of salient contextual features. I believe that the prospects for virtue contextualism finally depend to a considerable degree on spelling out a convincing account of the relevant salient features.

ACKNOWLEDGMENTS

I thank John Greco for his comments on an earlier draft and in conversation and an anonymous referee of *Erkenntnis* for his stimulating remarks on my discussion of the lottery problem.

NOTES

- ¹ In the penultimate version of Greco's paper, condition (c) had been formulated in terms of reliable cognitive character. I am puzzled why he changed it to cognitive abilities in the final version, because now it is hard to draw a clear distinction between conditions (b) and (c) and so his whole theory of intellectual credit seems somewhat unstable.
- ² An anonymous referee of *Erkenntnis* actually raised this objection and tried to support it with the following reason.

REFERENCES

- Cohen, S.: 1988, 'How to Be a Fallibilist', *Philosophical Perspectives* **2**, 91–123.
- Feinberg, J.: 1970, *Doing and Deserving – Essays in the Theory of Responsibility*, Princeton University Press, Princeton.
- Greco, J.: 2003a, 'Knowledge as Credit for True Belief', in DePaul, M. and Zagzebski, L. (eds), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Oxford University Press, Oxford, 111–134.
- Greco, J.: 2003b, 'Further Thoughts on Agent Reliabilism: Replies to Cohen, Geivett, Kvanvig, and Schmitt and Lahroodi', *Philosophy and Phenomenological Research* **66**, 466–480.
- Greco, J.: 2004, 'A Different Sort of Contextualism', *Erkenntnis* **61**, 383–400.
- Kvanvig, J.: 2003, *The Value of Knowledge and the Pursuit of Understanding*, Cambridge University Press, Cambridge.

Freie Universität Berlin
 Institut für Philosophie
 Habelschwerdter Allee 30
 14169 Berlin
 Germany
 E-mail: dirkkoppelberg@aol.com

PETER BAUMANN

LOTTERIES AND CONTEXTS

ABSTRACT. There are many ordinary propositions we think we know. Almost every ordinary proposition entails some “lottery proposition” which we think we do not know but to which we assign a high probability of being true (for instance: “I will never be a multi-millionaire” entails “I will not win this lottery”). How is this possible – given that some closure principle is true? This problem, also known as “the Lottery puzzle”, has recently provoked a lot of discussion. In this paper I discuss one of the most promising answers to the problem: Stewart Cohen’s contextualist solution, which is based on ideas about the salience of chances of error. After presenting some objections to it I sketch an alternative solution which is still contextualist in spirit.

1.

Here is a puzzle about knowledge. Stewart will never be a multi-millionaire (p), and he knows it. He also knows that if he will never be a multi-millionaire, then the lottery ticket he has just bought won’t win ($p \rightarrow q$); winning the lottery would make him a multi-millionaire. However, it seems that Stewart does not know that he won’t win the lottery (q).¹

We can call this problem “the Lottery puzzle”.² It constitutes a puzzle because we have several extremely plausible propositions here which are logically incompatible (with each other). People like Stewart (and most of us) can certainly know that they will never be multi-millionaires. Given our assumptions about Stewart it is thus very hard to deny that

$$(1) \quad Kp.$$

It is also hard to see how he could not know that not becoming a millionaire entails not winning the lottery. Hence, we may also assume that

$$(2) \quad K(p \rightarrow q).$$

Moreover, there is a well known principle of closure according to which knowledge is closed under known entailment:

$$(C) \quad [Kp \ \& \ K(p \rightarrow q)] \rightarrow Kq.$$

Not many would be willing to give up this principle.³ Now, from (1), (2) and (C) we can finally infer that



(3) Kq .

However, almost everybody seems to agree that nobody can know the outcome of a lottery in advance or just on the basis of the statistical evidence (even if that evidence is overwhelmingly in favor of a certain outcome). So (3) seems false. But doesn't it follow from true premises?

It seems that we must give up one of our propositions – but which one and why? If one does not give up the plausible closure principle, then we seem to have the choice between dogmatism and skepticism: Either we must, per *modus ponens*, accept that we know things we did not think we knew (3) or we must, per *modus tollens*, deny that we know things we thought we knew ((1) and (2)). Since many if not all ordinary propositions that we claim to know empirically (like p) entail some “lottery proposition” (like q), the puzzle can be easily generalized.⁴ It constitutes a general problem about knowledge.⁵

There has been a substantial amount of discussion in recent years about possible solutions to this puzzle.⁶ In this paper I will first discuss one of the most interesting proposals: Stewart Cohen's contextualist solution, which makes crucial use of the notion of salience of chances of error (I).⁷ After that, I will give a sketch of an alternative solution to the puzzle (II).

2.

According to Cohen,⁸ there is an implicit context-shift in the above inference from (1), (2), and (C) to (3). The standards of knowledge that allow for a true attribution of knowledge to Stewart that he won't be a multi-millionaire are not very high. In other words, the attributor's context is not a very demanding one.⁹ However, as soon as we start to think about the lottery, the probabilistic chances of error become salient. Moreover, the salience of the chances of error pushes the standards of knowledge so high up that it turns out to be false to say of Stewart that he knows that he won't win the lottery. The attributor now finds herself in a much more demanding context.¹⁰ All this allows Cohen to preserve closure: Given the same context (standards and salience relations), knowledge is closed under known entailment. If we change contexts during the inference, closure does, of course, not hold. But who would expect that? Cohen's solution to the lottery puzzle is apparently one more example of a contextualist solution to a gripping philosophical problem.

However, I think this solution is not quite convincing. It is *not* the salience of the chances of error that makes the difference and solves the problem. First, some remarks on salience.¹¹ Cohen holds that the growing salience of the chance of error can raise the standards for knowledge *so much* that knowledge is “lost” in the process (“lost” in the sense that it becomes false to say that the person knows that *p*).¹² To be sure, Cohen does not argue that the raise of standards always “destroys” knowledge. But he claims that in the Lottery case rising standards do make knowledge unattainable for us.¹³ This implies something that very often goes unnoticed: Salience is a matter of degree. Things are more or less salient to a person. And that which becomes salient – the chance of error – also admits of degrees. Hence, it is plausible to assume that both factors determine how much the standards of knowledge rise: The more salient a chance of error is and the greater the chance of error, the higher the standards of knowledge.¹⁴ Compare two different scenarios. In the first one, *S* plays a lottery with odds of 1 – 10. Suddenly the attributor becomes very much impressed (and has reason to become impressed) by the chance of error. In the second scenario, the odds are 1 – 2 million; the attributor is less acute and not too much impressed (and has no reason to become impressed) by the chance of error when it’s pointed out to him. Let us, for the sake of simplicity, assume that there is no further difference in (relevant) contextual factors and that potential losses and gains are the same in both lotteries. Wouldn’t we say that the standards for knowledge are much higher in the first scenario than in the second? Don’t we need a graded account of the rising of standards? So far, I think, Cohen can agree.

The crucial point, however, is the following one: Why should we believe that in a lottery the growing salience of the chance of error pushes the standards of knowledge *so* high up that knowledge disappears?¹⁵ Why should we assume that it becomes false to say the person knows that *p* (again, given that there are no blocking mechanisms at work here)? Consider a rather relaxed and not so acute attributor and ridiculously unfavorable odds; let us assume that Jack’s winning of the lottery is less likely than the end of the world by entropy within the next 5 min. Who would want to say that under these conditions the mere fact that the chance of error becomes salient pushes the standards for knowledge up *to the extent that* knowledge about the outcome of the lottery is “lost”? I do not see any reason to assume that. On the contrary: Sometimes we do know that we won’t win the lottery – even though the chance of error is clearly salient (I put this point aside here but will come back to it

later). Cohen's argument about salience does not seem to solve the lottery puzzle. In other words: There are lotteries and salience relations such that knowing that one won't win cannot be lost by simply making the chance of error salient – even *very* salient – and thus raising the standards. The crucial question here is not so much “Does the chance of error become salient?” but rather “Is that sufficient for losing knowledge?”. If knowledge can be lost in some but not all of those cases, in which standards rise and the chances of error become salient, then we need a different explanation. It might still refer to contextual factors, yet not to salience or to salience alone. Salience of error is not sufficient for loss of knowledge. I do not even see why it should be considered necessary.

Consider the well-known newspaper example.¹⁶ Suppose *S* participates in a lottery with odds of 1–1 million (and it later turns out that *S* does not win). Let us assume that the chance of error becomes salient (enough) for us, and that we judge that *S* does not know that he won't win. Later *S* reads the lottery results in the newspaper; according to the newspaper report, *S* has lost. We know that this newspaper reports incorrect lottery results in 1 out of 1000 cases. The chance of error is 1000 times greater here than in the first case. We can even be very much aware of the limited reliability of his newspaper. However, at the same time we still would not doubt that *S* acquires knowledge that his ticket has lost by reading the papers – even if we are as focused on the probabilities as in the lottery case. So, it cannot be reference to the salience of the chance of error that solves the Lottery puzzle.¹⁷ Cohen might disagree with this and say that in the newspaper case above we would not claim that the person knows she has lost – and would even claim that the person does not know she has lost – and that this is because we focus on the probabilities and the chances of error. However, I must say that I do not find this reply very plausible: Even very critical readers can learn something from the papers, and even the strongest fallibilist can accept that people come to know things from a source which is less than 100% reliable.

One last remark on Cohen's positive proposal. Cohen subscribes to fallibilism: the view that “I can know that *P* even though the probability of *P* on my evidence is less than 1”.¹⁸ How is this fallibilism compatible with his story about salience, that is, the view that knowledge attributions are false if made in the light of possible error? Can we only make true knowledge attributions – at least in “fallible” cases in which the probability of “*p*” on the evidence is less than 1 – while we are not aware of the truth of philosophical or commonsensical fallibilism? Because otherwise awareness of the fallibility in the

particular case at hand would make the chances of error salient and thus “destroy” knowledge? Is fallibilism true but not assertible or not believable – because we cannot think or talk in a fallibilist way about knowledge, given the salience principle? Contextualists like Stewart Cohen or David Lewis¹⁹ often say that when we think about skeptical scenarios we have to deny ordinary knowledge claims. Apparently, according to what Cohen says, we do not need to move into such extravagant contexts; thinking about fallibility is sufficient for making us deny ordinary knowledge claims (given that they refer to fallible cases). This is interesting because people do think about fallibility in perfectly ordinary contexts; hence, even in certain ordinary contexts, knowledge claims would have to be denied. It seems to me that this is a somewhat unwelcome implication of Cohen’s position.

3.

So, what should we say about the Lottery puzzle? I agree that we should not give up closure too easily.²⁰ I also accept Cohen’s claim that at least in some contexts it is true to say that he knows he will never be a multi-millionaire. The crucial question then is this one: Does he know that he won’t win the lottery? If not, why not?²¹ Why do we want to deny that Cohen knows that he won’t win the lottery? And what implications does all that have for closure?

We can get closer to an answer to these questions if we take an interesting principle about knowledge into account. This principle can, I think, help us explain our intuitions that we cannot know lottery propositions. Consider an ordinary case of knowledge: *S* knows there is still milk in the fridge and *T* does not (*T* just dogmatically assumes it). *S* knows it because he has looked inside the fridge whereas *T* doesn’t because he has not checked. There is a difference as to the quality of the epistemic position of *S* and *T* with respect to the proposition that there is still milk in the fridge. The difference between a good (enough) and a bad (enough) epistemic position explains why one person knows and the other does not. This gives us a principle that I propose to use for a solution of the Lottery puzzle:

- (EP) If *S* knows that *p*, then there are both good and bad epistemic positions for *S* with respect to the proposition that *p*.²²

For instance, *S* can have better or worse evidence for “*p*”. *S* can either shortly glimpse into the fridge or take a very close look inside

and check whether it really is milk rather than, say, yoghurt. *S*'s tongue might be better or worse at detecting whether "that stuff" really is milk or rather some kind of fake-milk (tofu-milk, etc.). If *S* relies on testimony that there is still milk in the fridge, his sources might vary as to their reliability, and *S* might be doing a good or not so good job at selecting and evaluating his sources. Insofar as deductive or probabilistic reasoning is involved, *S* might show different degrees of sophistication or lack thereof. Generally speaking, how good or bad the epistemic position of a person is with respect to some proposition depends on the information the person has, on her cognitive capacities and on the use she makes of those capacities.^{23,24}

Now, in Lottery cases all epistemic positions are "created equally".²⁵ In a lottery there is no difference between good and bad epistemic positions. To be sure: One person might have very vague or incorrect ideas about lotteries, another person might have illusions about the odds and think that winning the lottery is as probable as rain in Scotland tomorrow. However, we are obviously not dealing with such differences here. Rather, we are dealing with a subject who knows pretty well what a lottery is, is aware of the odds, etc. When we ask whether *S* knows that his ticket is a loser, we must presuppose that *S* is aware of the crucial facts concerning the lottery; otherwise it would be much less than clear whether we can ascribe beliefs about losing the lottery to the subject at all. Knowing as well as merely believing a proposition involves a sufficient grasp and understanding of that proposition.²⁶

Given all that, it seems very plausible to say that there is no difference between good and bad epistemic positions in lottery cases. How could there be such a difference? No matter what the person does or does not do – she cannot improve or spoil her epistemic position with regard to a particular lottery proposition. One could say (if it were not a bit too misleading) that the person plays an epistemic lottery. No matter what the person does, in a monetary as well as in an epistemic lottery, her winning or losing the prize (money, truth) does not depend on what she does. This is why we say that the person does not know she will lose the lottery.²⁷ She does not meet (EP). To be sure: According to (EP) the reason why one does not know that ticket no. 367 will lose is not that one's epistemic position with respect to this proposition is not different from one's epistemic position with respect to propositions about the other tickets; some authors have proposed such a view.²⁸ (EP) only has to do with different possible epistemic positions with respect to *one* such proposition.

(EP) explains why we think that one cannot know a lottery proposition. Apart from that, it throws some light on our concept of knowledge in general. A further advantage of this account is that (EP) also explains the similarities and dissimilarities between the Lottery and the Preface paradox.²⁹ In both cases, we have to deal with n items (tickets, beliefs), and in both cases, only one (or very few) of the items has an interesting property F (wins, is false); finally, the subject does not know which item has F but knows that some item has it. We may further assume that each of the true beliefs in the Preface scenario qualify as knowledge.³⁰ Hence, in the case of the Preface the following is true: In each case of a true belief the subject knows that the item does not have the interesting property (falsity). The case of the lottery is different: At least in some contexts it is not true to say of the subject that she knows that the relevant item (losing ticket) does not have the interesting property (wins). (EP) can explain this asymmetry. In the case of the lottery, there are no better or worse epistemic positions with respect to particular propositions like "Ticket no. 367 won't win". There is nothing I can do to improve (or spoil) the evidence I have for this proposition. In the case of the Preface, however, there is, according to the story, such a difference: The author has better or worse evidence for each of the individual propositions. For instance, if I have written a book on Italian cuisine, then I can have better or worse evidence for particular propositions like "The Italians had pasta already before Marco Polo went to China".³¹ It does not matter whether the epistemic positions are the same or similar for each of the propositions in the Preface case; what matters is that with respect to each individual proposition, I can be in a better or worse epistemic position (e.g., have better or worse evidence). It also does not matter if and to what degree the person in the Preface case is aware of the chances of error; even for a neurotic fallibilist who can never *not* think about his own fallibility, there will be this difference between the Lottery and the Preface. The Preface does not involve an epistemic lottery but the Lottery does. I think that a condition of adequacy for a solution of this Lottery puzzle is that it explains the similarities and dissimilarities between the Lottery and the Preface (a similar adequacy condition would hold for solutions of the Preface paradox). Cohen's proposal does not seem to meet this condition.³²

(EP) helps us to explain why we are inclined to say that the person does not know that she won't win the lottery. However, this is not yet the whole picture. At the same time, we often feel entitled to say that somebody can or does in fact know that she won't win the lottery.

There is something between knowledge in the above sense (involving (EP)) and complete ignorance. Consider the following dialogue:

- A: *Perhaps we should call off our plans for the summer.*
 B: *Why that?*
 A: *Well, I've bought lottery tickets. We may be multi-millionaires by Friday.*
 B: *C'mon – you know that you won't win the lottery!*³³

It seems to me that here lies the real context dependency: Sometimes we use looser standards of knowledge and then it is true to say that *S* can or does know he won't win the lottery; in these cases meeting (EP) is not necessary for knowledge.³⁴ In other contexts, we use stricter standards and add the additional conditions captured by (EP).³⁵ Then it is true to say that *S* cannot or does not know that he won't win the lottery.³⁶ There are, in other words, weaker and stronger standards for knowledge and they vary with context.³⁷ Sometimes, especially when it is necessary that we act soon or when we feel like acting soon or when we have the practical implications for everyday life in mind, we are happy with the weaker standards; however, when we have more “theoretical” interests, we might be more inclined to insist on conditions like those mentioned in (EP). I am not trying to explain here why the difference between good and bad epistemic positions matters in some contexts but not in others; I am just describing the phenomena. It seems pretty clear that all this has nothing to do with the salience of chances of error. Let us use “knowledge-l” for talk about knowledge at looser standards and “knowledge*” for knowledge at higher standards. The latter requires fulfillment of the condition mentioned in (EP): that it is possible for the subject to be in better or in worse epistemic positions with respect to the relevant proposition.

What are the implications of all this for the Lottery puzzle? It seems true to say of Stewart that he can know that he will never be a multi-millionaire. He knows it, given weaker as well as stricter standards (of knowledge-l or of knowledge*); there certainly is a difference between good and bad epistemic positions with respect to the proposition that Stewart will never be a multi-millionaire (he has recently checked his bank account, thought about different possible investments, discussed the future of the stock market, etc.; it is obvious that he could also not have bothered to think about money at all). But how is that compatible with the fact that in some contexts it is false to say that Stewart knows he will not win the lottery? Doesn't this violate closure?

Yes and no. Let me explain. When we are talking about knowledge in some weak sense (knowledge-I) in which it does not require a difference between good and bad epistemic positions, then closure holds: Stewart knows-I that he will never be a multi-millionaire, and he also knows-I that he won't win the lottery. However, if we switch to knowledge* closure does not seem to hold any more: Stewart knows* that he will never be a multi-millionaire (given the assumptions I made about him in the last paragraph), and he also knows* that this implies that he won't win the lottery; however, he cannot know* that he won't win the lottery. He can only know-I it ("know-without-a-star"). In some (stricter) contexts, closure does not hold, in other (weaker) contexts it holds. So, there is some consolation for the friends of closure. One can still be a friend of closure, if I'm right here – though perhaps not as close a friend as expected. But would that be so bad? And why not "contextualize closure", especially if you're a contextualist anyway?³⁸

We should not be surprised at all that closure does not hold in all cases, that is, for all kinds of knowledge or for all kinds of contexts and standards.³⁹ Take, for instance, the obvious case of different sources of knowledge and consider the following example about knowledge by testimony. I know by testimony both that my car has a problem with the cooling system, and that if one has such problems one will have to pay a lot for having it fixed. My mechanic just told me that. I infer from this that I will have to pay a lot for having it fixed. Even before my mechanic tells me and presents me with the bill, I know it will be expensive. Hence I know it by deduction but I do not know it by testimony. In other words, knowledge by testimony is not closed under known entailment. (Knowledge-I by inference, in contrast, always seems to satisfy closure).

The Lottery puzzle is puzzling insofar as it is not always clear to us that there are different contexts in play here. In this respect I am not that far away from Cohen. He is right in locating the problem in the context dependence of standards. But he is, I think, not right in bringing in the salience of chances of error. What really matters is rather our intentions, practical purposes and other contextual factors.⁴⁰ They determine whether weaker or stronger standards are adequate.

ACKNOWLEDGMENTS

For discussions and comments I am grateful to Stewart Cohen as well as to Bob Plant, Jon Cameron, John Greco, Christoph Jäger, Ryan

Nichols, an anonymous referee, and the audience at the Conference “Contextualist Approaches to Epistemology” (September 2003) at the University of Mainz (Germany).

NOTES

- ¹ Cf. Cohen (forthcoming). Cf. for an earlier presentation and discussion of the puzzle: Harman (1973, p. 161).
- ² This should, of course, not be confused with Kyburg’s lottery paradox (cf. Kyburg, 1961, pp. 197f.).
- ³ Strictly speaking, there is not just *one* closure principle but many different ones. For a very useful discussion of several options cf. Hales (1995). Here is an objection to (C): One might know that *p* and also know that it entails *q* but just not draw the inference. Closer to the truth is thus the following version: If one knows that *p*, and if one comes to accept that *q* on the basis of a correct inference, then one knows that *q*. But even this won’t do: Suppose Jack knows the grocery store is open. He has (properly) ignored the possibility that an earthquake might just have destroyed the grocery store but that does not undermine the knowledge claim. However, Jack cannot “bootstrap” and gain new knowledge that there has not been an earthquake by inferring this from the proposition that the grocery store is open. For a version of a closure principle that is closer to the truth cf., e.g., Wright (2000), Davies (1998) and Barke (2004). For the sake of simplicity, I will stick with (C). Nothing essential hinges on it and the same things hold *mutatis mutandis* for more adequate closure principles. I will also restrict myself here to single-premise closure and not go into multiple-premise closure (cf. e.g., Hawthorne, 2004, pp. 46–50).
- ⁴ A lottery proposition in the broader sense of the word would be any proposition that we think is most probably true but which we think we do not know (given the statistical nature of the evidence we have for it). For instance: I think that there is a very small probability that I will be hit by lightning tomorrow but do I know that I won’t be hit by lightning tomorrow? A lottery proposition need not be about a lottery. For more about examples like the latter one, cf. Vogel (1990). For the sake of simplicity, I will mainly deal with lottery propositions in the narrow sense of a proposition about the outcome of a lottery.
- ⁵ As Cohen (forthcoming) notes.
- ⁶ Cf., e.g., Olen (1977, pp. 521–523), Stemmer (1982), Adler (1986, pp. 244–248), Vogel (1990, p. 16, 22), Dudman (1992, p. 205), DeRose (1996), Ryan (1996, p. 130), Nelkin (2000, pp. 388–390), Weintraub (2001), Olin (2003, pp. 98–104), Greco (2003, this issue, and Hawthorne (2004).
- ⁷ Cf. Cohen (forthcoming, 1988, pp. 106–108; 1998). Similar strategies have been proposed by David Lewis and, to some degree, John Hawthorne: cf. Lewis (1983, 1999, pp. 421, 430, 443f), and Hawthorne (2002) (but see also Hawthorne 2004). For a different contextualist solution of the lottery puzzle cf., e.g., DeRose (1996).
- ⁸ Cf. Cohen (forthcoming). There are other important issues in Cohen’s article which I cannot discuss here. Main parts of it deal with Hawthorne (2004).

- ⁹ Cohen assumes attributor contextualism; I think he is right to do so.
- ¹⁰ Cohen seems to assume that every sentence expressing a lottery-proposition uniquely determines just one standard or context. I would rather take it that each such sentence comes with different standards, according to varying contexts. I will make use of this further “contextualization” below.
- ¹¹ “Salience” should not be taken in a purely psychological sense; salience does not reduce to what the person does in fact pay attention to but also involves what she should pay attention to.
- ¹² For the sake of ease of expression, I will from time to time use expressions like “loss of knowledge” rather than expressions like “change of the truth-value of ‘S knows that *p*’ from *True* to *False*”. Nothing substantial depends on this way of talking and there is no use-mention confusion going on here.
- ¹³ Cf. Cohen (forthcoming). It is presupposed here that there are no blocking mechanisms at work, like “C’mon manoeuvres”. A *C’mon* manoeuvre is an attempt to resist the change of standards, in particular their raise. Jack might claim to know that the grocery store is open. Jill might try to raise the standards by pointing out Jack does not know that no earthquake has just destroyed the store. Jack can try to resist this move by replying “C’mon, we both know that nothing like that has happened!”.
- ¹⁴ Things become more complex if we factor in potential gains and losses. The more pressing our need to get things right, the higher (*ceteris paribus*) the standards. One might suspect that the following comes close to the truth: The more salient a chance of error and the greater the chance of error as well as the potential gains and losses, the higher the standards of knowledge will be. For the sake of simplicity we may stick with the first principle above. It may not be more than a rule of thumb but it works at least in many cases.
- ¹⁵ Cf. also Greco (2003, this issue) on this point.
- ¹⁶ Cf. Cohen (1998, pp. 292f), Harman (1968, p. 166; 1986, p. 21).
- ¹⁷ DeRose (1996, pp. 576–579) makes a similar objection against Cohen. However, I do not find DeRose’s own solution convincing either.
- ¹⁸ Cf. Cohen (forthcoming, p. 2).
- ¹⁹ Cf. Cohen (1988, p. 105) and Lewis (1999).
- ²⁰ Cf., however, Harman and Sherman (forthcoming).
- ²¹ One might suspect that the person does not know that she has lost because she would still believe she has lost even if she had in fact won the lottery. Given that knowledge requires “sensitive belief” (cf. Nozick, 1981, p. 172ff.) she does not know she has lost the lottery. However, sensitivity accounts of knowledge seem to create more problems than they solve. I cannot go into this here (neither into other alternative accounts of knowledge).
- ²² (EP) might only hold for empirical propositions but that would be sufficient for my purposes here. – I will later have to restrict (EP) but for now we can leave it like that.
- ²³ There is a lot more to say about the notion of an epistemic position. In this context, however, the remarks above should be sufficient.
- ²⁴ Virtue epistemologists (cf., e.g., Greco, 2003, this issue) would probably replace talk about good and bad epistemic positions by talk about epistemic virtues and vices. What I have in mind here is quite different from virtue epistemology. First, my use of the terms “information”, “capacities” and their “use”, is purely descriptive, not normative. Second, informational inputs or the uses of certain

cognitive capacities need not have anything to do with overarching dispositions and stable character traits. Finally, nothing in the account proposed here suggests that one should explain knowledge and justification in terms of epistemic virtues – rather than the other way around, i.e., explain epistemic virtues in terms of knowledge and justification. – Some of the points John Greco makes come quite close to what I am saying here. However, in the end Greco says that it is the salience of chance in lottery cases that undermines crediting the true belief to intellectual abilities and epistemic virtues (cf. Greco, 2003, p. 8). This differs, of course, from what I am proposing here.

- ²⁵ Or, at least they are more or less equal. It does not matter here whether some tickets have a better chance of winning than others or whether there are, unknown by the ticket holders, no winning tickets in this particular run of the lottery at all; we can disregard these complications here for the sake of simplicity.
- ²⁶ This understanding need, of course, not be perfect. But it must meet minimal requirements. – This is not to deny that there are cases in which someone knows that some proposition he doesn't understand is true. I can, for instance, know by testimony that some scientific statement is true without understanding it. It seems obvious to me that we are not dealing with such cases here.
- ²⁷ One can read Adler (1986, p. 248), Olin (2003, pp. 98–104) and Greco (2003, this issue) as going into a similar (but not quite the same) direction.
- ²⁸ Cf. Harman (1968, pp. 166–168, 1973, p. 160f), Dretske (1981, p. 99), Adler (1986, p. 247), Vogel (1990, p. 22), Ryan (1996, p. 130) and Hawthorne (2004, 15f). I cannot go into a discussion of this alternative proposal here; I only mention that it seems that it cannot account for the similarities and dissimilarities between the Lottery puzzle and the Preface paradox (see below on this adequacy condition for any account of the Lottery puzzle).
- ²⁹ An author of a book has good inductive reasons to believe (and to say in the preface) that not everything he says in the book is true. It would be crazy and irrational to assume one's own infallibility. On the other hand, a sincere author has good reasons to believe every single thing he says in the book; hence, the author also seems to have good reasons to believe that everything in the book is true. This, however, is incompatible with the fallibilism mentioned at the outset. This paradox is not just one for authors but holds for all kinds of sets of beliefs. Cf. Makinson (1965). I will not consider the case here in which all of the beliefs in the book/the relevant set of beliefs are in fact true. I will also assume that the person knows that at least one of her beliefs is false.
- ³⁰ The fact that there is some doubt as to which belief might turn out false, does not speak against that; otherwise the Preface paradox would be a skeptical puzzle.
- ³¹ I once found this claim in a preface to a cook book; the author was Italian. The claim seems historically correct.
- ³² Neither do the accounts mentioned in fn. 28.
- ³³ On such *C'mon* manoeuvres see also Cohen (2004, p. 11), DeRose (forthcoming) and Hawthorne (2004, pp. 84, 161).
- ³⁴ I must leave it open here what knowledge requires in such weaker contexts: reliable true belief or something else or perhaps little more than true belief. It should be clear enough for my purposes here what I have in mind when I talk about knowledge in the "weak" sense.
- ³⁵ We now have to restrict (EP) to the latter kinds of contexts (I will not explicitly mention this from now on).

- ³⁶ I have my doubts that we can always rank standards according to strength but for the sake of the argument I will go with this assumption here.
- ³⁷ I am not even sure that there is a context-dependency here at all; perhaps the word “knowledge” is ambiguous but I won’t pursue this possibility here.
- ³⁸ The question here and in current debates on closure more generally is not whether principle (C) above or no closure principle at all is true. Rather, the controversial and difficult question is whether we need to restrict, modify and supplement (C) and if yes, in what ways. See also Hales (1995) here.
- ³⁹ Cf. also Bogdan (1985) who argues in a similar vein.
- ⁴⁰ Cf. Harman (1973, pp. 160f). – I acknowledge that more has to be said about these factors and how they work – but not here.

REFERENCES

- Adler, J. E.: 1986, ‘Knowing, Betting and Cohering’, *Philosophical Topics* **14**, 243–257.
- Barke, A.: 2004, ‘Epistemic Contextualism’, *Erkenntnis* **61**, 353–373.
- Bogdan, R. J.: 1985, ‘Cognition and Epistemic Closure’, *American Philosophical Quarterly* **22**, 55–63.
- Cohen, S.: 1988, ‘How to Be a Fallibilist’, *Philosophical Perspectives* **2**, 91–123.
- Cohen, S.: 1998, ‘Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery’, *Australasian Journal of Philosophy* **76**, 289–306.
- Cohen, S.: forthcoming, ‘Knowledge, Speaker, and Subject’.
- Davies, M.: 1998, ‘Externalism, Architecturalism and Epistemic Warrant’, in Wright C., Smith B. and Macdonald C. (eds.), *Knowing Our Own Minds*, Clarendon, Oxford, 321–361.
- DeRose, K.: 1996, ‘Knowledge, Assertion and Lotteries’, *Australasian Journal of Philosophy* **74**, 568–580.
- DeRose, K.: forthcoming, ‘Single Scoreboard Semantics’, *Philosophical Studies*.
- Dretske, F.: 1981, *Knowledge and the Flow of Information*, MIT Press, Cambridge, MA.
- Dudman, V. H.: 1992, ‘Probability and Assertion’, *Analysis* **52**, 204–211.
- Greco, J.: 2003, ‘Knowledge as Credit for True Belief’, in DePaul M. and Zagzebski L. (eds.), *Intellectual Virtue – Perspectives from Ethics and Epistemology*, Oxford University Press, Oxford.
- Greco, J.: 2004, ‘A Different Sort of Contextualism’, *Erkenntnis* **61**, 383–400.
- Hales, S. D.: 1995, ‘Epistemic Closure Principles’, *Southern Journal of Philosophy* **33**, 185–201.
- Harman, G.: 1968, ‘Knowledge, Inference, and Explanation’, *American Philosophical Quarterly* **5**, 164–173.
- Harman, G.: 1973, *Thought*, Princeton University Press, Princeton.
- Harman, G.: 1986, *Change in View. Principles of Reasoning*, MIT Press, Cambridge, MA.
- Harman, G. and B. Sherman: forthcoming, ‘Knowledge, Assumptions, Lotteries’, *Philosophical Issues*.
- Hawthorne, J.: 2002, ‘Lewis, the Lottery, and the Preface’, *Analysis* **62**, 242–251.
- Hawthorne, J.: 2004, *Knowledge and Lotteries*, Clarendon, Oxford.

- Kyburg, H. E. Jr.: 1961, *Probability and the Logic of Rational Belief*, Wesleyan University Press, Middletown, CT.
- Lewis, D.: 1983, 'Scorekeeping in a Language Game', in D. Lewis (ed.), *Philosophical Papers*, Oxford University Press, Oxford, Vol. 1, 233–249.
- Lewis, D.: 1999, 'Elusive Knowledge', in Lewis D. (ed.), *Papers in Metaphysics and Epistemology*, Cambridge University Press, Cambridge, 418–445.
- Makinson, D. C.: 1965, 'The Paradox of the Preface', *Analysis* **25**, 205–207.
- Nelkin, D. K.: 2000, 'The Lottery Paradox, Knowledge, and Rationality', *Philosophical Review* **109**, 373–409.
- Nozick, R.: 1981, *Philosophical Explanations*, Harvard University Press, Cambridge, MA.
- Olen, J.: 1977, 'Knowledge, Probability, and Nomic Connections', *Southern Journal of Philosophy* **15**, 521–526.
- Olin, D.: 2003, *Paradox*, Acumen, Chesham Bucks.
- Ryan, S.: 1996, 'The Epistemic Virtue of Consistency', *Synthese* **109**, 121–140.
- Stemmer, N.: 1982, 'A Solution to the Lottery Paradox', *Synthese* **51**, 339–353.
- Vogel, J.: 1990, 'Are there Counterexamples to the Closure Principle?', in Roth M. D. and Ross G. (eds.), *Doubting. Contemporary Perspectives on Skepticism*, Kluwer, Dordrecht, 13–27.
- Weintraub, R.: 2001, 'The Lottery: A Paradox Regained and Resolved', *Synthese* **129**, 439–449.
- Wright, C.: 2000, 'Cogency and Question-Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof', *Philosophical Perspectives* **10**, 140–163.

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

STEWART COHEN

REPLY TO BAUMANN

Peter Baumann raises several important issues regarding my proposal for resolving the lottery paradox.¹ The first concerns the role of salience in my account. According to my view, *S* can know *P* on the basis of evidence *E*, even though there is a possibility *H* compatible with *E*, but incompatible with *P*. That is to say, I endorse a form of Fallibilism. We must endorse some form of Fallibilism in order to avoid skepticism. All the same, it seems to be a fact about us that when error possibilities are salient to us, i.e. possibilities compatible with our evidence *E* but incompatible with *P*, we will deny, or at least feel a strong intuitive pull toward denying that *S* knows *P* on the basis of *E*. I attempt to explain this phenomenon by arguing that ascriptions of knowledge are context-sensitive. The standards for how strong one's evidence must be in order to know can vary with the context of ascription. So although in ordinary contexts, I can truly say "I know that I'll never get rich," in contexts where error possibilities are salient to me, the standards rise so as to falsify my knowledge ascription. In such a context, I cannot truly say, "I know I'll never get rich."

One way error possibilities can become salient to me is when my reasons are explicitly probabilistic. So for example, in a fair lottery with *n* tickets, my basis for thinking that my ticket will lose is the $n - 1/n$ probability that it will lose. But this basis for thinking that I'll lose makes salient the $1/n$ chance that I'll win. Thus, in typical contexts, the standards will be such that I fail to know my ticket will lose. And this is confirmed by the fact that the intuition that I don't know I'll lose the lottery, simply on the basis of the statistical probabilities is very robust.

Baumann raises some interesting questions about the role of salience in my account. Noting that salience comes in degrees, he argues that this is something my theory should take into account. The more salient an error possibility is, as well as the greater the chance of error, the higher the standards for knowledge should rise. I'm not sure if this is correct and I find the example he uses to motivate his claim somewhat unclear. He compares two lotteries, one where *S*'s



odds of winning are 1 in 10, the other where the odds of winning are 1 in 2 million. In the first case, the chance of error is very salient, whereas, in the second case the chance of error (with respect to *S* believing he'll lose) is only somewhat salient.² He suggests that we would say that the standards for knowledge are higher in the first scenario.

But I myself have no clear intuitions about the standards for knowledge in these cases. I do have intuition about who knows and who doesn't, and my intuition is that in neither case does *S* know he'll lose. So I do not see support for Baumann's view that degree to which the standards go up depends on the extent to which error possibilities are salient. On my view, when error possibilities are salient enough, the standards rise to whatever extent is required to falsify the knowledge ascription. So the degree of salience is relevant only insofar as error possibilities must be salient enough to motivate the intuition that the subject fails to know.

Baumann challenges my view that when the chance of error is salient standards rise to the point where knowledge ascriptions are falsified. He argues that if the odds of winning the lottery are high enough, then we are comfortable with saying that the subject can know he will lose the lottery. I find that in general, people have the intuition that one does not know one will lose the lottery merely on the basis of statistical probabilities, regardless of the number of tickets in the lottery. Baumann argues that where the number of tickets is so great that the chance of my winning is less than the chance that the world will end by entropy in the next five minutes, we will readily say that I can know that my ticket will lose. Thus salience of error possibilities is not sufficient to falsify knowledge ascriptions. I agree that this may be true. Certainly the intuition that I do not know my ticket will lose is much less strong in this case. But perhaps we can explain this by noting that we do not have a very clear conception of what it would mean for the world to end by entropy in the next 5 min. Thus the error possibility is too unclear to be sufficiently salient. Notice that most people will say that one fails to know one's ticket loses regardless of the number of tickets.

Baumann also says that he does not see why salience of error should be necessary for the loss of knowledge. Of course he is right here, but I have never claimed otherwise. One can lose one's knowledge simply by forgetting, or by getting new evidence.

Next Baumann turns to the newspaper example. Although we deny that I can know my ticket loses simply based on the number of tickets in the lottery, we readily allow that I could know my ticket

loses by reading about it in the newspaper. This is so, even though the probability that the ticket will lose based on the number of tickets can be greater than the probability that my ticket loses, based solely on the newspaper report. I have argued that this case supports Contextualism since it must be that the standards for knowledge are greater in the case where I believe my ticket will lose based solely on the number of tickets than in the case where I believe my ticket will lose based on the newspaper report. Why else would we allow that I can know in the case where the chance of error is greater? I also argue that this phenomenon comports well with my salience account. As I noted earlier, one cannot avoid thinking about the $1/n$ probability that one will win if one bases one's belief on the fact that the odds of losing are $n - 1/n$. So the chance of error will be salient in this case. But matters are different in case where I base my belief on the newspaper report. In such a case, we typically do not infer that what the newspaper says is true, based on the m/n probability that if the newspaper says it, it is true. Rather, insofar as we think of it at all, we think of a scenario: The reporter determines the result, perhaps by witnessing the drawing, sends in the report and then it is printed in the newspaper. In this scenario, the chance of error is not salient. Now if we did explicitly consider the ways in which newspapers can make mistakes, viz. newspapers do make mistakes, e.g., misprints, do we know that this particular report was not a misprint?, etc., we would feel the pull of saying that we did not know the newspaper report is true.

Baumann does not find this plausible. He notes that

Even very critical readers can learn something from the papers, and even the strongest fallibilist can accept that people come to know things from a source which is less than 100% reliable. (Bauman, 2004, p. 418)

But this point does not count against my view. I allow that we learn things from newspapers and I endorse the view, in the abstract, that people can know from a source that is less than perfectly reliable. More specifically I argue that we can make true knowledge ascriptions of this kind in everyday contexts where the less than perfect reliability is not salient. But I hold that in particular instances where the chance of error becomes salient, the standards rise so as to falsify these particular knowledge ascriptions.

The same point applies to Baumann's claim that my account of how standards get raised is incompatible with my endorsement of Fallibilism. How can I at one and the same time think that knowledge ascriptions are fallible but true, if I think that when such fallibility is

salient, we will deny that we know? Again, the point is that we endorse in the abstract that our knowledge ascriptions are fallible and yet sometimes true. It is only when error possibilities are made salient with respect to a particular case, that we become reluctant to affirm that we know in that case.

Baumann proposes his own account of our lottery intuitions. The reason I fail to know my ticket will lose, in a fair lottery, is that one's epistemic position in a lottery case is fixed. There is no prospect of one's position getting better or worse. Baumann holds that this conflicts with what he takes to be an axiom about knowledge:

- (EP) If *S* knows that *P*, then there are both good and bad epistemic positions for *S* with respect to *P*.

But even if we grant the principle, I do not see how my knowing I lose the lottery would violate EP. Let's suppose that I cannot improve my situation. The odds are what they are and there is nothing else I can do. Still, I am in a rather good epistemic position with respect to the proposition that my ticket is a loser. Someone who has absolutely no idea what the odds are, is in a bad epistemic position with respect to that same proposition. And someone who has already witnessed the drawing, supposing it has already taken place, is in a much better position than I am. My not being able to improve my position is quite a different matter from there not being good and bad epistemic positions.

But perhaps Baumann means that there must be good and bad positions that I myself could occupy. Even granting this interpretation of the principle, I do not see how the lottery case fails to meet the principle. Suppose the drawing has already occurred but I still have not heard the results. I could certainly be in a better position in the case were I to have witnessed the drawing.

Baumann also argues that my account is inadequate because it fails to explain the difference between the lottery paradox and the paradox of the preface. But the difference can be explained by noting that in the lottery, we are considering a specific proposition, *My ticket is a loser* and we can very easily see how we could be wrong about this, viz., my ticket could be drawn. In the preface there is only the abstract worry that when there are so many propositions, the probability that one is false is high. But this is not the same thing as it being salient with respect to a particular proposition how I could be wrong about it.

Baumann's own proposal for explaining these paradoxes has the result that deductive closure for knowledge is falsified, even relative

to a context. On my view, that is a very serious difficulty for a view, since deductive closure seems compelling. Baumann tries to mitigate this result by noting that deductive closure fails for particular kinds of knowledge. I can know P by testimony, deductively infer Q and thereby come to know Q . But I do not know Q by testimony. Therefore I know P by testimony, but I fail to know a deductive consequence Q by testimony.

But I do not think that this kind of closure principle for kinds of knowledge has anything like the intuitive plausibility of deductive closure *simpliciter*. Certainly when we come to know things by perception or testimony, there will be deductive consequences that we do not know by perception or testimony. This is a far cry from allowing that I could know that I'll never get rich yet fail to know that I won't win the lottery, even though winning the lottery will make me rich.

NOTES

¹ Baumann, "Lotteries and Contexts", *Erkenntnis* **61**, 415–427.

² Baumann talks as if salience is a matter of how impressed one is or how much reason one has to be impressed. But on my view, salience is purely a psychological notion akin to attention.

Department of Philosophy
 Arizona State University
 P.O. Box 874102
 Tempe, AZ 85287-4102
 USA
 E-mail: cohen@asu.edu

MARK LANCE and MARGARET LITTLE

DEFEASIBILITY AND THE NORMATIVE GRASP OF CONTEXT

ABSTRACT. In this article, we present an analysis of defeasible generalizations – generalizations which are essentially exception-laden, yet genuinely explanatory – in terms of various notions of privileged conditions. We argue that any plausible epistemology must make essential use of defeasible generalizations so understood. We also consider the epistemic significance of the sort of understanding of context that is required for understanding of explanatory defeasible generalizations on any topic.

1. INTRODUCTION

There are many generalizations in epistemology which, though full of exceptions, seem essential to the very enterprise of knowledge. “The future will be like the past,” “things are as they appear to be;” “people tell the truth;” “well established methodological principles can safely be employed barring specific reason for doubt.” In none of these cases is the generalization universal or exceptionless. There are any number of cases in which appearances are deceiving, the future is not like the past, and methodologies which, though long-established, are systematically misleading. Nonetheless, deployment of such generalizations is essential in epistemic life; to abandon them would leave us little idea how to proceed in the game of rationality.¹

Indeed, it is not uncommon for epistemologists and philosophers of science to suggest that all the interesting epistemic generalizations are like this. Fallibilism, epistemic holism, Neurath’s boat, Quine’s web, Sellars’s diachronic process – all seem to have as a consequence the idea that what counts as evidence is intrinsically contextual. A consideration that in one situation counts for a conclusion can, in another, count against it: if the experience as of seeing a cup sometimes counts as evidence that there is, indeed, such a cup, there are other contexts – say, when one has just ingested an hallucinogen, when it counts against it; and it’s doubtful we could spell out in any finite, concrete terms the cases in which it does and doesn’t. Non-trivial justification is always defeasible.²



In this regard, contextualism in epistemology seems analogous to claims made in ethics by a view known as “moral particularism”. Moral particularism is, very roughly, the view that considerations’ reason-giving force is irreducibly context-dependent.³ Considerations, it is said, carry their moral import only holistically: a consideration that in one context counts for an action can, in another, count against it or be irrelevant, and all in a way that cannot be cashed out in finite helpful terms. The claim is not simply that the moral contribution made by these considerations can get *outweighed* by others (as when the pain of a shot is justified by the utility it brings), but that the moral “valence” of the consideration, as it were, itself changes. Pain is bad – well, except when it is constitutive of athletic challenge; intentionally telling a falsehood is *prima facie* wrong – well, not when done to Death Squad agents, to whom the truth is not owed, or when playing the game Diplomacy.⁴

But comparisons to moral particularism should give us pause. If claims of contextualism are widely regarded as uncontroversial in epistemology, the claims of moral particularism are widely regarded as controversial at best, downright crazy at worst. Many reasons for worry have been raised. First, it is said, moral particularism seems to imply that there is no more intimate a connection between honesty and the good than between, say, shoe-lace color and the good. After all, each can, in the right context, be good- or bad-making – or neutral; but it seems odd to think that honesty is only accidentally related to moral status. Second, it is claimed, the view seems to imply that there is no structure to moral theory at all. Moral understanding must be simply a matter of accumulating a series of one-off pieces of insight (*x* is good here; *y* is good there; do *z* next Thursday in Pittsburgh), a picture which makes it puzzling how morality could be learned, debated, improved upon, or even discussed.

If these problems are serious in ethics, they should be for epistemology as well. If the epistemic “valence” of appearances, for instance, can switch, how do we recover the idea that there is some important theoretical connection between having an appearance that *p* and being justified in believing that *p*? If there is no explanatory or in some way law-like structure to be found, then how would the standards of epistemic justification be any easier to learn, debate, etc. than standards of moral good?

For both disciplines, strategies for answering these queries divide into two broad camps. “Radicals” embrace – indeed, celebrate – the claimed absence of law-like structures: understanding the morality of a given action, or again the justification of a given belief, it is said, is a

matter of *discerning* how the moral or epistemic considerations add up in each individual case. Understanding is not a matter of applying theoretical generalizations, for there are no non-trivial ones; it is, rather, a matter of skill or wisdom – the practice of moral and epistemic virtues. “Moderates,” in contrast, argue that this misreads the lesson of contextualism. The presence of exceptions does not mean we have left theoretical generalizations behind, but that we must reconceive what those generalizations are like if they are to do justice to the nature of these realms. More specifically, we must soften their semantic content and recognize their nature as “*ceteris paribus*,” “hedged,” or “defeasible” generalizations: defeasibly, pain is bad-making; defeasibly, appearances are to be trusted.

We believe the second camp has to be the right one: a discipline – be it ethics or epistemology – empty of any theoretical or law-like generalizations is a discipline with highly attenuated potential for understanding. But the usual proponents of this strategy, we want to argue, end up giving interpretations of defeasible generalizations that abandon the central contextualist insight. More specifically, we will argue, they either bank on the possibility of expunging exception at one level or another of theory, or they leach away the ability of the theoretical generalizations to do the explanatory heavy lifting we expect any such claims to do. At the deepest level, we believe, such interpretations still believe that exceptions are at war with explanation.

In what follows, we give a different account. We argue that there are defeasible generalizations whose semantics must be understood in terms of a normative conception of “privileged conditions.” This approach provides us with a crucial tool for understanding the role of explanatory exception-laden generalizations in epistemology (and ethics); in addition, it helps us to understand how we could come to understand generalizations that are robust while full of holes. Reflection on their nature, in short, enables us to make progress both on the role of generalizations in epistemology and on the epistemology of generalizations.

2. DEFEASIBLE GENERALIZATIONS

The qualifier “for the most part” gets bandied about rather casually, as do its rough surrogates “*ceteris paribus*,” “all things equal,” “as a rule,” “defeasibly,” etc. These phrases are often used interchangeably and without much explanation. However opaque their use, they are ubiquitous:

- Defeasibly, matches light when struck.
- *Ceteris paribus*, lying is wrong-making.
- For the most part, pain has a negative valence.
- Other things being equal, fish eggs develop into fish.
- In standard condition, red ties look red.
- As a rule, the future is like the past.
- Subject to provisos, an increase in supply leads to a drop in price.
- Generally, people say what they believe.

What *do* we mean when we advance a generalization thusly qualified? Attempts to explain the family of qualifiers fall into two broad traditions. On the first tradition, such claims are read as purely statistical ones. On this interpretation, “for the most part” means, quite literally, that the asserted connection holds with high frequency. They are statements about what is likely to happen in one’s own neck of the woods, and can be used to underwrite, amongst other things, default presuppositions – assignments of non-inferential entitlements or “start here” positions.⁵

To read such generalizations in this way, however, is to give up any pretense that they are robustly explanatory. Except in areas like quantum mechanics, which are ruled by genuinely statistical laws, statistical generalizations are contingent ones. Moreover, the interpretation is just factually unavailable with many central examples, such as the claim about fish eggs – which, as it turns out, only rarely succeed in turning into fish.

The second tradition preserves the explanatory nature of “for the most part” generalizations by interpreting that qualification as a signal that we are talking shorthand. The assertion is seen as an enthymeme – a claim containing suppressed premises. There is a concrete exceptionless generalization in the offing, we just don’t quite know yet what fills in the gaps or it’s so obvious we needn’t bother to state it. On this reading, “for the most part” generalizations are indeed explanatory, but they are thought capable of serving that function only because the exceptions are in principle eliminable.

In its straightforward form, this approach has fallen into universal disfavor in epistemology. The idea seems to be of a piece with early positivist attempts to understand explanation in terms of deduction from exceptionless principles. While officially rejecting hypothetico-deductivist approaches, however, it turns out that much recent work on *ceteris paribus* generalizations retains essential allegiance to the enthymemist’s conviction that the route to genuine explanation involves finding some way to eliminate exceptions.

Thus one strategy states that CP clauses serve to restrict the *scope* of the explanatory generalization. On this view, CP generalizations are exceptionless within a certain boundary and cases outside that boundary are *terra incognita*.⁶ Of course, we do not deny that there are explanatory generalizations which are thus restricted. But notice that the systematic application of this strategy is just another version of the enthymematic approach: we can translate “CP, all As are Bs,” into “All As within the boundary are Bs.” More fundamentally, such an approach seems unable to account for the full power such generalizations can have, for boundary-specific universal generalizations are altogether silent on what is going on outside the boundary. Applied to our generalization about appearances, this approach would interpret us as saying merely that appearances within the boundary are veridical – leaving open the possibility that appearances become utterly irrelevant, as opposed to complexly relevant, outside that boundary. We, however, are confident that the CP relevance of appearance to truth is of more systematic epistemic importance than this.

On another approach to defeasibility, the CP operator is a signal that we are isolating a constant force. Thus, if Newton’s laws were meant to be descriptive laws of motion – e.g. an object in motion stays in motion with the same vector – they would require a CP clause, since obviously objects don’t always move in a straight line. (Similarly, one cannot simply read the CP here as “in the absence of other competing forces.” This restricted boundary reading makes the generalization come out true, but since it says nothing about any circumstance in which there *are* other forces – i.e. any real circumstance – it robs the resulting generalization of its explanatory power.) The move, then is to say that the laws are capturing, not regularity of actual motion, but rather an underlying force: it is always the case that the force of inertia is equal to the vector of motion, but this is one determinant of motion among many.

Philosophers such as John Earman insist that genuine sciences must always expunge CP clauses laws in some such manner. Earman actually goes further to insist that, like Newtonian force analysis, respectable laws must have available a combinatorial algorithm to take us from all such forces to a net result of motion. The first part of this move – the move to forces – is represented in ethics by the view of W.D. Ross, but Ross resists the latter move. Moral generalizations capture constant moral “forces” – they isolate features that are always “good-making” or “bad-making” – but there is no algorithmic way to combine these. Reaching an overall conclusion on an act

requires, according to Ross, moral judgment, which is an inherently skill-like form of perception. In the philosophy of science, Nancy Cartwright presents a view rather like Ross's. She urges that explanatory generalizations all have to do with capacities or powers – elements which are uni-directional – but not algorithmically combinable in all cases. (Cartwright also believes that there is a strong contextualist qualification in most actual scientific explanations, indicating a combination of the delimitation and the force strategy.)⁷

Here again, we do not deny that CP clauses can function to isolate forces; we deny, though, that they always so function. For, crucially, there are important explanatory generalizations in epistemology, as in ethics, in which the very “valence” of the explanatory concepts can change with context. We discuss such cases in detail in the next section. For now we note merely that the “force” interpretation of CP clauses shares in the presumption that explanatory work is ultimately carried out only by exceptionless generalizations. Our primary goal is to show that this need not be assumed.

If we leave aside the specifics of the force idea, we see that the strategy of relocating explanation to an area expunged of exception is extraordinarily common. In ethics, for example, many philosophers have argued that, when confronted with a moral principle inherently involving a CP qualifier, one might have a genuine explanation, but only if that explanation is *underwritten* in one way or another by a *non-defeasible* generalization at a more abstract level. An action's honesty can explain its rightness, even allowing that in other contexts it wouldn't, but, these abstractionists would urge, this is only because there is a more fundamental law, say, “fidelity is *always* good-making, and honesty only sometimes constitutes fidelity.” Whatever the details, such views suggest that we need to abstract from the immediate phenomena to find the fundamental level at which exceptionless laws govern. Once again, we find the fundamental premise that real explanation happens only where exception no longer resides.⁸

We find little reason to find this view plausible in epistemology. What would the more abstract concept be which governed generalizations involving appearances? Perhaps some candidate could be proposed, but we doubt it; in any event, our goal will be to show that one needn't search for such an underlying level of theory. Exceptionless generalizations, we will be arguing, are perfectly respectable in their own right.

While both the statistical strategy, and all the various versions of the enthymematic strategy are often useful, they don't exhaust the possibilities. We want to suggest another interpretation of *ceteris*

paribus or hedged generalizations, one in which the contextualist means to capture something quite different. On this interpretation, such qualifiers are used to point to a kind of generalization that is both genuinely explanatory *and* ineliminably exception-laden – and as it will turn out, irreducibly normative.

3. DEFEASIBILITY AND PRIVILEGED CONDITIONS

When we issue a generalization to the effect that something has a certain feature, sometimes what we really want to say is not that such a connection always, or even usually, holds, but that the conditions in which it *does* hold are particularly revealing of that item's nature. We might put it by saying that we are asserting what happens in “normal” conditions, except that the notion of “normalcy” is so freighted with misleading connotations. Better put, then, we are taking as *privileged*, in one way or another, cases in which the item has the feature specified. Such generalizations can tell us about the nature of something, then, not by eliminating exceptions to the connection, but by maintaining and demarcating their status as exceptions.

Take a plebian example. Defeasibly, soccer (“football” as the rest of the world insists on calling it) is played with 11 members on a team. Only defeasibly, for there are any number of variations – pick-up soccer with three on three, “little league” soccer with 20 on 20 and no goalie, the list goes on. There is surely no specifying out in any concrete terms when a game counts and when it doesn't. (We can codify – in excruciatingly boring detail – the structure of FIFA soccer; but there is no codifying which pick-up games count as a riff rather than a different game.) The variants, indeed, probably statistically predominate. Nonetheless, it would be quite wrong to think the play with 11-members is just one among many. The other games are, crucially, understood by reference to the first. Games may count as soccer while deviating from this standard, but it is nonetheless the standard. 11-member soccer stands in no need of explanation, while other versions do and the explanations reside in appropriate relations to the standard case.

Further, coordinate revisions in the rules are justified in terms of the ways the number of players deviates. In 5 on 5 soccer, one typically has a smaller goal. Why? Because otherwise, given the smaller number of players and consequent increase in open space, there would be too many goals – that is, enough to constitute an unacceptable deviation *from standard soccer*. That is, it is not that there exists some Platonic norm opposing games with lots of goals (one

needn't eschew basketball to motivate a smaller goal in soccer). Rather, one sort of soccer is functioning as a standard, and acceptable variations are motivated by their relation to this norm. In this sense, then, even non-standard soccer games carry a "trace": they each defeasibly involve 11-membered teams, in the sense that their deviations from 11 membered soccer must be justified and shown to be acceptable as variations.

In this case, we are not using the *ceteris paribus* generalization to say that soccer usually has the features highlighted. Nor, though, must we think we can exhaustively specify the conditions under which it in fact would. One needn't specify the conditions in which a connection does obtain in order to say that where it does it counts as privileged and privileged in a way that allows the generalization to explain why we make certain other changes.

Suitably developed, this view allows us to understand defeasible generalizations in a way that retains their utility for epistemic theory. In overview, when we say "as a rule, appearances are to be trusted," or again "*ceteris paribus*, the future will be like the past," we are neither saying that these regularities always hold, nor merely asserting that they usually do so in our parochial locale; we are endorsing, instead, defeasible generalizations, committing ourselves to the claim that the conditions in which they do hold are epistemically privileged. Where appearances, past experience, or accepted practice are not-trustworthy, it is in virtue of the ways they deviate from privileged conditions, and one does not adequately understand the relevant terrain who does not appreciate the nature and import of that deviance.

Of course, to say without further analysis that qualified generalizations involve privileging moves would just add one more unexplained gesture to the pantheon. The task is to say something informative about privileging. Crucially, we want to argue, there are a number of distinct kinds of privileging moves: if defeasible generalizations are sometimes unclear, it is in part because the privileging moves undergirding them are complex. Before moving to a (partial) typology of the sorts of moves we have in mind, however, we pause to discuss a bit of the logical structure of defeasible generalizations.

4. THE LOGICAL GRAMMAR OF DEFEASIBILITY CLAIMS

Defeasibility, or so we allege, requires us to add just one operator to the pre-existing logical tool-box, but not an operator standing for "it is defeasibly the case that." Rather, we employ a modal operator *P*,

read simply as “in privileged conditions” While the structure of various porous generalizations will vary, we can make use of the usual conditionals, quantifiers, and connectives along with P to render a wide range of such claims.

Formally, the P operator will simply be a modal operator the semantics of which is given by designating a range of privileged worlds in a usual possible worlds framework. So $P(A)$ is true (at a world α) iff A is true at every privileged world (relative to α) In some areas of application the world relativisation of privilege will be important.)

To see why we can’t simply read “defeasibly A ” as $P(A)$, let us consider another simple example:

- (1) Defeasibly, matches light when struck.

A key feature of this generalization is its non-monotonicity. Defeasibly matches light when struck, but it is not the case that defeasibly, matches light when struck underwater. Indeed, wetness is a paradigmatic defeator of lighting and part of the grasp of the artifact kind “match” requires grasping defeasibly, matches struck underwater do not light. But if we read 1 simply as

- (1a) $P(\forall x \in m)(x \text{ is struck} \rightarrow x \text{ will light})$

it would entail

- (2a) $P(\forall x \in m)(x \text{ is struck underwater} \rightarrow x \text{ will light})$

In privileged conditions, after all, matches aren’t underwater. The whole point is that being under water is a *non-privileged* condition. It is not, however, *altogether* under-privileged: non-monotonicity expresses itself precisely in the fact that there are even less privileged conditions under which matches underwater do light. Rather, it seems that we need to look at the most privileged underwater conditions, that is, conditions understood as differing from the privileged ones “as little as possible,” that is, in only those ways regarded, *from the point of view of the privileged world*, as going naturally with being underwater.

An important point: we should not think of this as some neutral notion of “least difference;” rather, what we want is a substantive sense of “normal” deviation. There is no non-perspectival sense of similarity in which a world with 5 on 5 soccer and normal size goals is less similar to a FIFA world than one with 5 on 5 and smaller goals. Quite the contrary. Just as a substantive part of our understanding of soccer is an appreciation of the fact that FIFA soccer is the privileged version, so too is our understanding of the relevant similarity

relation, our understanding that similarity in goal *difficulty* is more important than similarity in goal *size*. But this understanding is a substantive understanding of soccer, not a matter of general logical or semantic competence. An understanding of privilege and, hence, an understanding of defeasible generalizations and the concepts they govern, requires, then, an ability to view whole regions of possibility space from the perspective of the privileged worlds.

The suggestion, then, should be clear. What is wanted is a reading of (1) as

- (1b) $P(\forall x \in m)(x \text{ is struck} > x \text{ lights})$ where
 $>$ is the standard subjunctive conditional.

That is, in all privileged conditions, matches are such that if they were to be struck, they would light. Or, in all the worlds nearest to the privileged ones in which matches are struck, they light. Similarly, we have:

- (2b) $P(\forall x \in m)(x \text{ is struck underwater} > x \text{ does not light})$

In any world among those nearest to the privileged one in which a match is struck underwater, it fails to light. The non-monotonicity of defeasibility, then, is understood in terms of the non-monotonicity of subjunctives.

But not subjunctives alone. We might ourselves live in watery Atlantis where matches have become useless. We want to claim, nonetheless, that the artifact concept match still lives (or dies) by its association with a set of (dry) privileged conditions under which it lights. That is, even if most matches don't light when struck, it is still true of every match that defeasibly, it would light if struck. Similarly, even if the actual world is watery-albeit-liquid-oxygen-infused-Atlantis, it is still true that defeasibly wet matches don't light.

The connection between striking and lighting, note, marks precisely the difference between seeing something as of the artifact kind "match" and seeing it as of the natural kind "magnesium-tipped stick." In any possible world, if there are matches in that world then in worlds which are privileged (with respect to that world, and for the concept "match") those matches light when struck. To be committed to this necessary defeasible generalization is part of what it is to treat something as a match – this, even if one occupies circumstances in which they rarely or never do light. Of course, which artifactual categories one bothers to maintain has an indirect dependency on such statistical frequencies: Atlantans presumably wouldn't bother demarcating something as this particular artifact. The point, though,

is that the artifact kind where deployed is circumscribed by marking some conditions – however frequent or rare they may contingently be – as privileged.

Though much more could be said about the formal structure of various sorts of defeasibility claims – especially around the very interesting ways that the modal structure interacts with quantification in universally quantified defeasible generalizations – we turn now to some clarification of the notion of privilege that we have been helping ourselves to, outlining two of the central-most such moves.

5. CLASSIFICATORY DEPENDENCE

Defeasible generalizations sometimes serve to classify something as the type of thing it is. The soccer and match examples, in fact, display this sort of reliance, but let us give another everyday instance to begin more detailed exploration. Most people would be willing to define a chair, functionally, as something to sit on. A moment's reflection, however, reminds us that there are any number of exceptions to this generalization – ornamental chairs made intentionally frail, for instance. Moreover, it is hard to see how we could say once and for all what counts – when the object at the Museum of Modern Art is a chair, or a work of art, or both.

Nonetheless, it seems right to think there is an intimate connection between the concept “chair” and the function of holding people in repose; and we might intuitively think to put the point by saying something like “*ceteris paribus*, chairs are things we can sit on.” Such a claim is not a statistical one. A very opulent, or orthopedically challenged, society might in fact have more ornamental than functional chairs lying about. Nor are we saying that there's anything defective about the exceptions (they are no good to sit on, to be sure, but that doesn't keep them from being fabulous – and fabulously sought-after – chairs). What we mean, instead, is that all chairs in the privileged class are fit for sitting, and that the relation between privileged and peripheral chairs is something like that of theme and variation. The ornamental chair is, if you like, a riff on the theme of chair; and one can't understand a riff without understanding the theme to which it stands as variation. (This is of course close to what Wittgenstein had in mind as a family resemblance: our point is that conceptual families, like their human counterparts, are rarely egalitarian.) The privileging move here, then, is about what has, as it were, conceptual priority: *to understand something as an ornamental chair*

one must understand the notion of chairs that are for sitting on, but not vice versa.

In this case we had a single concept, the extension of which was given via some notion of acceptable variation from a paradigmatic theme. Ornamental chairs are still chairs, though they are so by way of similarity to a different sort of chair. But paradigm-riff privileging often comes in a richer form. Consider irony. An ironic use of a sentence is a speech act in which what is meant is roughly the opposite of that which is usually meant by the utterance of that sentence. (“Well, George Bush sure proved a wise and peace loving leader.”) But irony is not simply a species of ambiguity, in which a sentence said in one tone of voice has one meaning and in another the opposite. For irony to function as it does, it must wear its reversal of semantic valence *on its sleeve*. It presents itself explicitly as being a non-standard use. Not, again, in a statistical sense: we could, in principle, turn into a society of Oscar Wildeses, using irony more than literal speech. The point is that these speech acts nonetheless function by *carrying a trace* of “standard” use. Utterances of *P*, we might put it, *always* have the property of defeasibly meaning *P*, even when used ironically to mean not-*P*. To use a sentence ironically is thus to use it in a way that can be understood only as derivative upon literal uses. Irony is essentially a riff on literal use, but a riff whose character as a riff is essential to it. (Compare kitsch in art.)

The category of “chair,” of course, is highly convention-dependent. For a case of classification via privileging that is less so (or at least less obviously so), we return to the fish egg example in biology. It is not only true of, but essential to the nature of, fish eggs that “other things equal, they develop into fish.” Here, as we mentioned, the statistical situation is even worse than usual: in many species, the vast majority of eggs never develop into fish, most being eaten or destroyed. Nonetheless, there is a crucial privileging of cases in which they do. One classifies the fish egg as being the kind of biological organism it is by reference to its “standard” or “normal” development. There are of course an infinite number of trajectories that fish eggs could take, from developing into fish, to being enucleated with sheep DNA and becoming a sheep, to breaking down into nutrients for a turtle, to being irradiated and turning into a strange and dysfunctional pile of flesh. Nonetheless, we elevate one such trajectory as a “natural” one, viz. one that does not call for explanation (at least at this level of theory); and in this sense, we circumscribe some developments as expressions of an organism’s “nature.” (Thus the fish egg is a *potential* fish, and a salamander egg – which could in some

possible world be turned into a fish by laboratory machinations – is only thereby a *possible* fish.)

In Aristotle’s day, of course, the “nature” of biological organisms was thought a practice-independent fact; in modern times, one may well think it a reflection of our classificatory scheme (which is not, of course, to say that it is shallowly conventional, without pragmatic constraint other than that there be a convention, as in which side of the road we drive on). The point is that to treat something as a biological nature is to elevate as unproblematic a given developmental trajectory. What it is to be a fish egg is to be the sort of thing that is in standard conditions when it develops into a fish and in non-standard conditions – conditions calling for explanation by way of riff on the standard – when it does not.

This notion illuminates well epistemic trust. Defeasibly, testimony of people (linguistic agents, epistemic agents) is reliable. Though there are any number of exceptions – both specific instances in which we have reason to doubt that a person is telling the truth, and also more general classes of cases such as topics that most people are unreliable concerning, as well as individual people who are frequently unreliable – it seems essential to the very category “person” that we see as privileged those cases in which their saying it justifies our (non-inferentially) believing it. As in the case of chairs, we can easily make sense of unreliable people, but to do so involves seeing them as non-privileged versions of reliable people, as like people except embodying some epistemic deviance – either a vice like dishonesty or gullibility, or unfortunate circumstances such as having been inserted into a nasty AI environment. Without reference to the usual case – without, that is, seeing the person as a riff on the theme “reliable epistemic agent” – there would be no grounds to treat an unreliable informant as speaking a language after all.

Note that this constraint is superficially similar to Davidson’s principle of charity. Indeed, in a number of places Davidson does express the principle as that there must be a defeasible rule to treat alien utterances as true.⁹ But he consistently goes on to give such statements a probabilistic reading – viz., it is a constraint on adequate interpretation that we interpret the majority of their claims as true. Now this statistical claim might state a necessary constraint on adequate interpretation (in fact we are far from convinced, but leave that aside). The central problem with the statistical principle of charity is that it isn’t *sufficient* to describe the operative interpretational constraint. The radical interpreter’s task is not merely to make sure that some sufficient aggregated majority – say, 51%, 75%, or 99%

of the interpreted utterances – are true. Rather, *in every case* in which interpretation attributes falsehood, the interpreter undertakes a particular responsibility to see that interpretive move as an acceptable riff on the theme of attributing reliability. *Every* deviation from the principle of charity wears its non-standardness on its sleeve; every deviation from the principle must be accounted for in light of the privileged relation to other people (for instance, “They get this wrong because they can’t see red” – because they differ in this understandable way from the normal person to whom we would attribute true beliefs).¹⁰

6. JUSTIFICATORY DEPENDENCE

A second central kind of privileging – though the categories are not exclusive – is justificatory dependence. To illustrate, imagine having a perception as of a red cup. Having such a perception – its appearing to one that there is a red cup, or as Sellars put it, “being appeared to red-cup-ly” – typically has a positive epistemic valence *vis a vis* the belief that there is a red cup; put into our language, defeasibly, appearances that *P* are justifying of beliefs that *P*. The claim is not that one typically does, or ought to, *infer from* the fact that it appears that there is a red cup to the *conclusion* that there is a red cup: the belief is non-inferential. Rather, the issue is what facts contribute to, or explain, the fact that one is justified – which facts are “justifying,” as we might put it, where “justifying” is used on epistemic analogy with the ethical “good-making.” Something that detracts from the justification of *P*, whether in the direction of justifying $\sim P$, or simply in the direction of reasonable doubt, we will call “skeptifying.”

There are, of course, contexts in which having such an appearance is skeptifying rather than justifying – as when you remember you have taken an hallucinogenic drug, or know the evil demon is playing with your eyesight in a particular way. Moreover, many will argue that there is no spelling out once and for all, in any relatively concrete terms, the conditions under which the perceptual experience is justifying.

Nonetheless, it seems natural to think there is some sort of intimate connection between appearance and justification. When appearances are unreliable – when seeing as *P*, or appearance that *P*, is not justifying of *P* – one’s knowledge of this fact itself relies on justification provided by contexts in which one can rely on appearances (as when, say, we see the evil demon at work). Cases in which

one is justified in taking one's appearances at their word stand as epistemically unproblematic; it is cases in which one is not so justified that demand explanation – and an explanation precisely that appeals to cases of the former type. Appearances, then, can mislead, but the relation between an appearance that *P* and a justified belief that *P* is deeper than the connection between, say, a justified belief that *P* and a justified belief that *Q* – even when *P* and *Q* happen to be tightly evidentially related; and this is so, even if given one's own background beliefs, the second actually holds more often in your vicinity than the first. For while the belief that *P* may in fact provide evidence that *Q*, it is of the essence of an appearance that *P* that it is defeasibly connected to justification of *P*. Appearances, we might put it, are necessarily *defeasibly* trustworthy. They carry this feature – the property of being defeasibly trustworthy – as a trace even into situations in which their justificatory import of trustworthiness changes from positive to negative.

Now amongst the exceptions we can encounter to being able to take appearances at face value, some are cases in which something is epistemically amiss: one is in a worse situation, by knowledge's own lights. Someone who is just entered the Hall of Holograms is in a situation which, however fun, is epistemically deficient. But exceptions do not always point to epistemic woe. If we have available a clear translation manual, we can make adjustments that preserve justification. Think of the well-medicated schizophrenic who continues to "hear voices as" telling her to do something – who continues, that is, to hear them non-inferentially as voices – but who is well trained to make an inferential adjustment to disbelief. Indeed, there are any number of cases in which we become quite practiced at inferentially adjusting what conclusion we draw when presented with an appearance: the bent stick is actually straight, the object looming in the rear view mirror is closer than it appears. Armed with the appropriate inferential adjustments, such situations need not be epistemically defective: one could know just as much in that situation as in the normal one. But the privileging relations nonetheless still apply. Situations in which appearances are taken at face value are privileged over equally epistemically productive situations in which they are not because our justification for our knowledge in the latter case is still dependent upon the former (we know that things appear differently in the mirror because, say, we observe without a mirror that they are so). Non-standard cases, thus, while always perform deviant, are not always defective.¹¹

When we say “all things equal, appearances are the sorts of things we can take at face value,” then, we are not claiming that appearances are *usually* trustworthy. Someone who gets stuck in the Hall of Holograms may never again be able to trust her eyes; the unfortunate brain in a vat is misled most of the time. What we are saying, rather – as an expression of a weak, non-foundationalist sort of empiricism – is that cases in which they *are* trustworthy are justificatorily basic in the sense that any knowledge, even the knowledge that one is in a situation in which one shouldn’t trust appearances, must rest on justified appearances. (Compare the Ptolemaic version of Plato’s principle “do it with circles,” in which any deviations from circular planetary orbits have to themselves be explained in terms of circles.)

This point helps to illustrate an important difference between something carrying a *defeasible import* and operating as an epistemic *default*. To call something a default is to say that it is a justified “start here” position – an assumption one is entitled to make in the absence of special evidence to the contrary. Now we believe that defaults are essential to epistemology; in particular, we believe that they are essential if we are to be protected from skeptical collapse. For all that, though, they are thoroughly distinct from defeasible connections. A defeasible import is the import something has in suitably privileged conditions, the understanding of which is thought crucial to understanding the theoretical significance of the kind. A default import is the import that an epistemic agent, in her local epistemic enclave, is justified in assuming a property has until presented with evidence to the contrary. A defeasible import *may* function as a default; but it precisely won’t when one is in a deviant situation.

Thus, in normal conditions, one is entitled to take it as a default that people are who they appear to be. When you see what appears to be your friend Jones, you get to assume without argument or question that it is indeed Jones, unless and until objections are raised. But conditions are not always normal. If you are at a costume party in which the norm of arriving in convincing costume is tightly binding, seeing that someone across the room appears to be your friend Jones may well in itself make it *less* plausible that the person is Jones. Not only, then, is trusting the appearance here *not* a default, the opposite is: Jones would never come looking like Jones, but Smith might think it a cute joke to dress as Jones. In this case, one wants to say that the relevant epistemic terrain is governed by the following norm: Defeasibly, one can treat as a default that people are who they appear to be.

Note now the relation between classificatory and various sorts of justificatory privileging. Certainly not all classificatory privileging is justificatory, but at least some justificatory privileging is classificatory, for justificatory privileging can serve to classify something as a kind. Sellars's view of perception is committed, in essence, to just such a claim. Sellars supposes that we believe, not only in episodes of being appeared to red-ly, but also in red appearances: While he maintained that appearances are not justificatory in abstraction from their being embedded in systems of beliefs about the world, that is, he famously argued that one must posit *sensa* as part of the explanatory story of how people come to be able to know by perception at all. If so, then it seems clear that, on Sellars's view, it is of the essence of the kind "appearance" that it be defeasibly justifying of the corresponding claim about the world. Just as with the difference between "match" and "phosphorus tipped stick," the crucial difference between a neurophysical category and the epistemic category of "appearance" is that the latter is "defeasibly to be believed." What is different in this case is that this conceptual dependency functions via a complex structure of justificatory and explanatory priority, a structure we have only scratched the surface of here.

7. RECOVERING THEORY IN A WORLD OF EXCEPTION

If the view being sketched here is right, then most epistemic generalizations, like most moral generalizations, are irreducibly porous. They are shot through with exceptions we cannot eliminate. These generalizations can nonetheless count as robustly explanatory and insightful.¹² Adducing them has a power a list of instances does not, for it situates instances within a framework that maintains some as exceptions to others' rule. Attempts to replace such generalizations with statistical or enthymematic surrogates result in a loss, rather than a gain, in explanatory power.

More specifically, epistemology can maintain a radical position on the valence-switching capacity of evidential considerations without saying that the hard-won lessons of philosophical reflection are merely useful local inductive tools, or pedagogic crutches to be left behind when enlightenment hits. Shoelace color and empirical appearance may each serve at times to justify a claim about the world, but the similarity ends there. Shoelace color doesn't have an epistemic nature; appearance does.

On the view being defended here, then, epistemology is far more than pure contextual know-how forever rooted in the particular. One must grasp principles to understand the very concepts of justification, knowledge, appearance, evidence, and even context itself. But it is also clear that in order to make sense of what it is to grasp such principles, “interpretation” will have to take on a more complex role in epistemic competence than even neo-Heideggerians have appreciated. The great majority of people working on the epistemology of principles these days – whether in morality, epistemology, or the philosophy of science – agree that competence requires interpretation of situations, rather than algorithmic application of principles stated in context neutral terms. When we turn to defeasible generalizations, though, the need for interpretation extends to the concept of privileged conditions itself. Interpretation arises in understanding what counts as a paradigm example, what counts as an acceptable deviation from that paradigm, and what follows from the way that an acceptable deviation deviates. Grasp of one’s situation requires grasp of one’s position *vis a vis* normalcy, a recognition of the trace left by the relevant defeasible generalizations.

This allows us to recover room – as surely we must – for theoretical understanding in epistemology. To be sure, if one’s only grasp of whether a condition was privileged *vis a vis* a given defeasible generalization was in terms of whether that generalization was there defeated, then defeasible generalizations would not function as genuinely explanatory. Such a view would involve treating “sets of privileged worlds” as mere formal book-keeping devices, perhaps useful for regimenting logical structure, but offering no independent epistemic leverage. If so, defeasible generalizations would merely make explicit our underlying skill and not function as genuine theory.

But there is no reason to think this the correct view. One of the things our accumulations of experience and developing epistemic skill begins to give us is a grasp of the shape of “privileged conditions.” If we are willing to countenance any sort of fundamental reliance on skillful understanding in epistemology, there is no reason it shouldn’t function here. To the extent we achieve such a grasp, it can inform what counts as an instance of a defeasible generalization. It is possible to learn about privilege – to see the shape of the set of privileged conditions – in a way that allows our conception of privilege to serve as a legitimate ground for revising our underlying inferential skills without that understanding – that sight – taking the form of an articulable theory. One can come to have a relatively independent partial understanding of the shape of the relevant concepts such as

normalcy, and even one that would serve as a legitimate ground for revising our underlying or antecedent inferential dispositions, without supposing that that understanding consists in possession of a tacit codifiable theory.

None of these issues of interpretation would even be visible to one not clear about the structure and importance of defeasibility. Barring the creation of an exhaustive exceptionless theory of privilege, navigating the world remains at bottom a matter of skill – including now a skill at understanding and recognizing what is deviant and normal, what paradigmatic and emendational, what conceptually prior or central. We must know our way around possibility space in a far richer sense than has previously been appreciated.

Compare a remark of Wittgenstein's:

But, after all, the game is supposed to be defined by the rules! So, if a rule of the game prescribes that the kings are to be used for drawing lots before a game of chess, then that is an essential part of the game. What objection might one make to this? That one does not see the point of this prescription, perhaps as one wouldn't see the point either of a rule by which each piece had to be turned round three times before one moved it. If we found this rule in a board-game we should be surprised and should speculate about the purpose of the rule. (*Philosophical Investigations*, §567)

In this general area of the *Investigations*, Wittgenstein was getting at two correlate errors. On the one hand, one could think that the game is either nothing but the rules, or again the actual practice of playing. (All commentators see that he is rejecting rules as definitive of practices; fewer realize he also rejects actual practice as thusly definitive.) On the other hand, one could think that there is some codifiable point to the game that could be used to define correct play. Neither is correct. Given the practice we find ourselves engaged in – and only from the perspective of some such engagement – we have a sense of the point of that practice, and understanding of our goals and purposes that allows us to amend that practice. But apart from our skillful involvement with the practice, we could not formulate any conception of its point, much less produce a codified theory of it that could be used to determine appropriateness within the practice.

A deep dependence on context, then, is consistent with the possibility of epistemic theory. Since the theoretical generalizations that comprise it involve a structure among situations of epistemic privilege and relations to privilege, though, theory looks very different from how it would look as a collection of exceptionless generalizations. Most deeply, a defeasible approach to epistemic theory argues that theory's usual quest, which is to spend all our time filling in the holes of our generalizations, is deeply misguided. We achieve wisdom,

in the end, not by filling in the exceptions, but by knowing what counts as one in the first place.

NOTES

- ¹ Of course none of these generalizations is uncontroversial. For purposes of this paper, all that matters is that there are some such generalizations which function as substantive explanatory generalizations in epistemology – generalizations, that is, which explain why a claim has the epistemic status it has – and which are, nonetheless exception-laden. What they are is of secondary importance. For present purposes we simply assume that there are some such generalizations, and take the ones mentioned to be plausible examples.
- ² The caveats “substantive” and “nontrivial” are, of course, crucial here. We can get obviously univalent epistemic considerations either by moving to the very abstract – positive reason to believe *P* makes *P* more justified – or to the very concrete – the observation that *P* at time *t* in situation *s* makes *Q* more justified.
- ³ See Mark Lance and Margaret Little, “Particularism and Anti-Theory.”
- ⁴ The first example is from Elijah Milgram; the third, from David McNaughton in conversation.
- ⁵ In ethics, this Millian strategy has been followed recently by Garrett Cullity, in “Particularism and Presumptive Reasons.” Few in epistemology have adopted this account of defeasible reasoning explicitly, but it has obvious resonances with reliabilism.
- ⁶ For discussion of this sort of approach, see Sylvain Bromberger, *On What We Know We Don’t Know*.
- ⁷ For a recent version of this strategy in ethics, see Paul Pietroski, “Prima Facie Obligations: Ceteris Paribus Laws in Moral Theory.”
- ⁸ See for example Roger Crisp, “Particularizing Particularism,” and David McNaughton and Piers Rawlings, “Unprincipled Ethics.”
- ⁹ Davidson himself does not distinguish adequately “defeasible” and “default;” see below.
- ¹⁰ Note that this suggestion puts attributions of falsehood in Davidsonian interpretation not merely into the camp of riffs on paradigmatic themes, but into the class of irony.
- ¹¹ A similar structure, we would claim, is exhibited in the case of the principle “the future will resemble the past.” Here the priority exhibited is primarily explanatory. One can, of course, justifiably posit any number of deviations from this norm – and it is not clear what sense could be given to counting how many – but the priority of the usual case of constancy seems to lie in the requirement that deviations be explained in terms of more basic constancies.
- ¹² We have not, of course, offered an account of explanation here. That is, we have not given any indication of when a defeasible generalization is explanatory and when it is not. Just as in the case of exceptionless generalizations, some are and some are not. The point is to offer an account of defeasible generalizations that allows one to adapt promising accounts of explanation to accommodate them as examples. Though it would take us rather far afield to survey the various theories, our claim is that isolating the core of defeasible generalizations as lying in a

universal claim about what happens in a privileged range of contexts, allows us to sidestep any of the standard arguments against the possibility of explanatory defeasible generalizations.

REFERENCES

- Aristotle: 1941. 'Nichomachean Ethics, and Posterior Analytics', in *The Basic Works of Aristotle*, Random House, New York.
- Bromberger, S.: 1992, *On What We Know We Don't Know: Explanation, Theory, Linguistics, and How Questions Shape Them*, University of Chicago Press, Chicago.
- Cartwright, N.: 1999, *The Dappled World*, Cambridge University Press, Cambridge.
- Copp, D. (ed.): forthcoming, *Oxford Handbook of Ethical Theory*, Oxford University Press, Oxford.
- Crisp, R.: 2000, 'Particularizing Particularism', in Hooker and Little.
- Cullity, G.: 2002, 'Particularism and Moral Theory: Particularism and Presumptive Reasons', *Aristotelian Society* 76 (Suppl.), 169–190.
- Dancy, J.: forthcoming, *Practical Reality*, Oxford University Press, Oxford.
- Davidson, D.: 1984, *Inquiries into Truth and Interpretation*, Oxford University Press, Oxford.
- Earman, J. and J. Roberts: 1999, 'Ceteris Paribus', There is No Problem of Provisos', *Synthese* 118(3), 49–78.
- Hooker, B. and M. Little (eds.): 2000, *Moral Particularism*, Oxford University Press, Oxford.
- Irwin, T. H.: 2000, 'Ethics as an Inexact Science: Aristotle's Ambitions for Moral Theory', in Hooker and Little.
- Lange, M.: 2000, *Natural Laws in Scientific Practice*, Oxford University Press, Oxford.
- Lance, M. and M. Little, forthcoming, 'Defending Moral Particularism', in J. Dwyer (ed), *Debates in Ethical Theory*, Cambridge University Press, Cambridge.
- Lance, M. and M. Little, forthcoming, 'Particularism and Anti-Theory', in *Oxford Handbook of Ethical Theory*, Oxford University Press, Oxford.
- McNaughton, D. and P. Rawlings: 2000, 'Unprincipled Ethics', in Hooker and Little, 256–275.
- Pietroski, P.: 1993, 'Prima Facie Obligations: Ceteris Paribus Laws in Moral Theory', *Ethics* 103(3), 489–515.
- Pietroski, P. and G. Rey: 1995, 'When Other Things Aren't Equal', *British Journal for the Philosophy of Science* 46(1), 81–110.
- Pietroski, P.: 1993, 'Ceteris Paribus Laws in Moral Theory', *Ethics* 103(3), 489–515.
- Ross, W. D.: 1930, *The Right and the Good*, Oxford University Press, Oxford.

Department of Philosophy
Georgetown University
New North
Washington, DC 20057
USA
E-mail: lancem@georgetown.edu; littlem@georgetown.edu

NIKOLA KOMPA

MORAL PARTICULARISM AND EPISTEMIC
CONTEXTUALISM:
COMMENTS ON LANCE AND LITTLE

ABSTRACT. Do we need defeasible generalizations in epistemology, generalizations that are genuinely explanatory yet ineliminably exception-laden? Do we need them to endow our epistemology with a substantial explanatory structure? Mark Lance and Margaret Little argue for the claim that we do. I will argue that we can just as well do without them – at least in epistemology. So in the paper, I am trying to very briefly sketch an alternative contextualist picture. More specifically, the claim will be that although an epistemic contextualist should commit himself to epistemic holism he can nevertheless appeal to epistemic principles other than defeasible generalizations in order to provide his epistemology with a structure.

1.

According to Brad Hooker, Moral Particularists “hold that the very same properties may count morally in favour in some circumstances and against in other circumstances” (Hooker, 2000, p. 6). Or in the words of John McDowell: “Occasion by occasion, one knows what to do, if one does, not by applying universal principles, but by being a certain kind of person: one who sees situations in a certain distinctive way” (McDowell, 1997, p. 162). Jonathan Dancy combines these points in his account of *Moral Particularism* (MP, for short). Dancy characterizes MP as “the claim that there are no defensible moral principles, and that moral thought does not consist in the application of moral principles to cases . . .” (Dancy, 2001, p. 1). And the core particularist doctrine, which Jonathan Dancy calls *the Holism of Reasons* (MH, for *Moral Holism*) is, “the doctrine that what is a reason in one case may not be a reason at all in another, or even a reason on the other side” (Dancy, 2001, 3, cf. also Dancy, 1993).

Dancy then goes on to distinguish two conceptions of moral principles: the absolute conception and the contributory conception.



According to the absolute conception, a moral principle is a universal claim to the effect that all actions of a certain type are overall wrong (or right). So there might, e.g., be a principle that says “be just.” But according to the contributory conception, “this does not mean that all just actions are in fact right; it only means that the justness of an action counts in its favour, or that an action is the better for being just” (*ibid.*, p. 6). The Particularist denies the existence of absolute principles. But he also denies the existence of contributory principles because he does not think it possible to specify the regular contribution a contributory principle would have to make (– if it is to deserve its name).

As Lance and Little note, there are serious problems for MP. Most importantly, it “seems to imply that there is no structure to moral theory at all” (Lance and Little, 2004, p. 436). More specifically, one problem is to explain how we can learn from our moral experience if we did not thereby learn to extract principles which we can avail ourselves of in new cases. A related problem is how we could learn the moral concepts in the first place. Suppose you just learnt that a certain concept *C* is applicable to cases *a* and *b*. What enables you to apply *C* to a new case, *c*, if not the following: you extracted the features which cases *a* and *b* had in common and you see that case *c* has these features too. You see, e.g., that the actions in case *a* and *b* were wrong because a promise was broken in each case. And in case *c* a promise was broken too. So you conclude that the action in case *c* was wrong too. But how did you do that if not by basing your judgment on your prior knowledge of a moral principle, namely that it is wrong to break a promise? This connects up with a third problem, the problem of how to justify our moral judgements. The particularist is not denying that there are moral reasons. Nor is he denying that an action’s being a promise-breaking might, in a particular case, be a reason for judging it wrong. But the fact that it was a promise-breaking in itself does not give the particularist any reason for judging the action wrong. It is exactly not the case, given his doctrine, that an action is wrong *just because* it is a promise-breaking. To try to explain that an action is wrong merely on the basis of the fact that it is a promise-breaking is simply no option for the particularist.

2.

Lance and Little claim that there is an analogy between MP and contextualism in epistemology. Moreover, they seem to suggest that

contextualism in epistemology is committed to a sort of epistemic holism. So let us characterize – in analogy to MP and MH – two epistemic positions: *Epistemic Particularism* (EP) and *Epistemic Holism of Reasons* (EH). Epistemic holism would be the view that what counts in one context for a given belief, might count in another context against it. Epistemic particularism, on the other hand, would be the claim that there are no defensible epistemic principles. Now Lance and Little further note that while MP is considered by many to be a very unattractive position, EH is rather uncontroversial. But if, as Lance and Little suggest, the one is analogous to the other, then shouldn't the problems for MP be equally problems for EH? It certainly seems so. But Lance and Little now argue that the problems can be solved if we acknowledge the existence of what they call defeasible generalizations.

There is an issue here of whether someone endorsing MH is thereby also committed to MP – and accordingly for EH and EP. I take it that the answer is no. Moreover, it seems as if the Moral Particularist faces the above-mentioned problems not because he is committing himself to Moral Holism, but because he is denying that there are any moral principles whatsoever. In any case, I will claim that we can commit ourselves to EH without thereby being forced to deny the existence of epistemic principles. EH can be separated from EP. So we may well accept EH and at the same time acknowledge epistemic principles other than defeasible generalizations. We should be willing to concede, though, that the principles might turn out to be very complex or even beyond detailed specification. But that is not to deny that there are any.

3.

Lance and Little elaborate on the notion of a defeasible generalization as follows. A defeasible generalization is a generalization that is genuinely explanatory but nevertheless ineliminably exception-laden. And, roughly, “Defeasibly, *P*” is to be read as “In privileged conditions, *P*”, where *P* is simply “a modal operator the semantics of which is given by designating a range of privileged worlds in a usual possible worlds framework” (*ibid.*, p. 443). So the question is: When are conditions privileged? It depends, for there are various distinct kinds of “privileging moves.” Lance and Little discuss two of them: classificatory dependence and justificatory dependence.

I won't have anything to say concerning classificatory dependence. It seems a fairly reasonable claim that we sometimes use defeasible generalizations in order to classify. So the idea of classificatory dependence and the conception of defeasible generalizations in general may well prove very useful. I am not disputing that. All I want to claim is that things are somewhat different when we turn to the epistemic case – i.e. to justificatory dependence.

So here is an example of a defeasible epistemic generalization: “Defeasibly, appearances that *P* are justifying of beliefs that *P*” (Ibd., p. 448). The point is that cases in which appearance is trustworthy are privileged. More specifically, the privileged cases exhibit a “justificatory and explanatory priority” (*ibid.*, p. 451). They do so because the cases in which appearance is deceptive call for an explanation which can be given only by appeal to cases in which appearance is trustworthy: The only way to find out that in a particular case appearance was deceptive is by using the senses, by reliance “on justification provided by contexts in which one can rely on appearances” (Ibd., p. 448). (Are there any other epistemic principles? The following might be suitable candidates: “In privileged conditions/defeasibly, testimony of other people is trustworthy.” “In privileged conditions/defeasibly, the deliverances of reason are trustworthy.”)

Now I don't think that we need defeasible generalization in epistemology (except maybe for classificatory purposes, but let us put that aside). In what follows, I will try to sketch an alternative contextualist picture. More specifically, I will claim that there are good reasons to endorse epistemic holism but that the tie between EH and EP should be severed. So while Lance and Little argue that the problems which EH inherits from EP can best be solved by appeal to defeasible generalizations, I will argue that the principles in question shouldn't take the form of defeasible generalizations if they are to prove useful in epistemology. For that purpose, I will focus on the following three questions: (1) Are we committed to Epistemic Holism? (2) Does Epistemic Holism generate problems, which can only be solved by appeal to defeasible generalizations? (3) Are we able to justify epistemic claims without having to avail ourselves of defeasible generalizations?

4.

Ad (1). The Epistemic Holist's basic claim is that a consideration can count for something in one case and against it in another case. In other words, the claim is that something can be a reason in one case

but no reason at all in another case. But what exactly does it mean to say that something is a reason in one case but no reason at all in another case? Suppose you look at a white wall which is, unbeknown to you, illuminated so as to look red. Now one might hold, along the particularist lines and in accordance with EH, that in this case, the fact that the wall looks red gives you no reason to believe that it is red. Alternatively, one could claim that it gives you a *prima facie* reason to believe that the wall is red – in analogy to W.D. Ross's Prima Facie Duties. An undercutting defeater can neutralize a *prima facie* reason, though. (For a definition of different kinds of defeaters cf. Pollock 1986, pp. 38/39.) This will be the case when you learn that there is red light installed in the room. Now which answer one prefers depends, among other things, on whether one thinks of epistemic reasons as being internally or externally constituted (or both). The Epistemic Holist seems to highlight the external aspect. In the Rosian picture, the focus is on the internal aspect of reasons.

But that is only part of the Epistemic Holist's story. And there might be a better reason why we should side with the epistemic holist on this score. Consider the following example of Crispin Wright's:

- (*e*) Jones has just headed the ball into the net, he is being
congratulated by team-mates and the crowd has gone wild.

This provides defeasible evidence for

- (*P*) Jones has just scored a goal,

which entails (assuming that it is only in the context of a soccer game that a soccer goal can be scored) that

- (*I*) A game of soccer is taking place.

But suppose the circumstances are special: you are in the vicinity of a film studio which specialises in making sporting movies and that you know that it is just as likely that the witnessed scene is specially staged for the camera as that it is an event in a genuine game. Once you are equipped with this information, you will rightly regard *e* as providing *no* warrant for *P*. What you need, if *e* is to provide warrant for *P*, is precisely some *independent* corroboration of the context – that is, of *I*. You ask a bystander: is it a genuine game or a film take? If you learn the game is genuine, you acquire warrant for the claim that a goal was scored. But it would be absurd to regard that warrant as transmissible across entailment from *P* to *I* (Wright, forthcoming).

Wright uses the example to illustrate the *information dependence of warrant*, which he characterizes further thus: "A body of evidence, *e*,

is an information-dependent warrant for a particular proposition *P* if whether *e* is correctly regarded as warranting *P* depends on what one has by way of collateral information” (Ibd.). And a warrant will fail to transmit “when the particular *e*, *P* and *I* have the feature that needed elements of the relevant *I* are themselves entailed by *P* (together perhaps with other warranted premises). In that case, any warrant supplied by *e* for *P* will not be transmissible to those elements of *I*” (Ibd.).

Now I take the claim that warrant is information-dependent in this sense to be a very reasonable claim. For the present purpose, though, I would prefer to put the point in terms of epistemic reasons instead of warrant. So we might put the point thus: sometimes evidence *e* gives you a reason to believe a proposition *P* only given some collateral information *I*. If *I* were false, then *e* would be no reason to believe *P*. *I* is to be presupposed if *e* is to give you a reason to believe *P*. So the important point is this: something is a reason for a given belief only against the background of certain other assumptions.

Suppose the wall looks red to you. This gives you a reason to believe that it is red – but only given that lighting conditions are normal. If that assumption is not in place, then the fact that the wall looks red gives you no reason to believe that it is red. And the point is not only that it gives you no “external reason” – i.e., no reason externally construed. It also gives you no “internal reason” because no evidence, taken in isolation, can function as a reason for anything. Something is a reason only within a network of other considerations. Nothing is a reason *simpliciter* – or so it seems in the light of the information dependence of warrant. So whether evidence *e* gives you a reason to believe *P* depends on what background assumptions are operative in the context. But then it seems that the information dependence of warrant commits us to EH (– something Lance and Little can agree with).

5.

Ad (2). But note that this is not a totally unrestricted anything-goes holism of reasons. It is not that anything can be a reason for anything else – without there being any theoretical structure. Let us take an example: *A* asks whether *B* has a reason to believe *P*. More specifically, the question is whether evidence *e* gives *B* a reason to believe *P* in the case at hand. Now, given the information dependence of warrant, it systematically depends on contextual features whether *e* is

a reason to believe *P*. It depends on what one has by way of background assumptions. But according to the contextualist version of EH I would like to put forth here, it depends not only on *B*'s background assumptions but it also on what background assumptions the attributor *A* makes. So there is a systematic connection between features of the subject's context and the attributor's context on the one hand and something 'a reason for a given belief in the context on the other. But the connection, though systematic, is apt to be highly complex. So contextualism is trying to reduce the complexity to a manageable size. And even if it would turn out that the connecting principles at work are so complex as to persistently defy specification, we could nevertheless try to approximate a specification.

Think of a context, very roughly, as of a set of propositions, comprising (i) the shared presuppositions – where these presuppositions need not be something which one has earned an independently established warrant for, (ii) the available collateral information – the assumptions for which an independent justification can be provided, and (iii) the set of assumptions under investigation. Obviously, this conception of context could be spelled out in terms of Robert Stalnaker's idea of a discourse context. Stalnaker proposes "to identify a context (at a particular point in a discourse) with the body of information that is presumed, at that point, to be common to the participants in the discourse" (Stalnaker, 1999, p. 98).

So again, *A* wonders whether evidence *e* gives *B* a reason to believe *P*. Now according to the contextualist, whether *e* gives *B* a reason to believe *P* depends on what is known, on what is presupposed, and on what is under consideration. But it depends not only on what is known or presupposed by *B*, but also on what is known or presupposed by the attributor *A* and on the purpose and direction of his inquiry.

Unfortunately, not all presuppositions are legitimate. So we need to be able to distinguish those, which are legitimate from those, which are not. And now someone might demur that it is here where the need for defeasible generalizations makes itself felt. For how else could we distinguish between legitimate and illegitimate presupposition? (How else, that is, could we meet the skeptical challenge). Is not a natural suggestion that in privileged conditions our presuppositions are legitimate, while in non-privileged conditions they are not legitimate? This brings us to the third question, the question of whether we need defeasible generalizations for justificatory reasons.

6.

Ad (3). Tom claims to know that there is a cat in front of him. John challenges the claim in the following familiar-sounding way: “How can you know that there is a cat in front of you? For all you know, you might be a Putnamian Brain in a Vat. And you can know that there is a cat in front of you only if you know that you are not a Brain in a Vat.” One way for Tom to try and meet the challenge is by pointing out that it appears to him as if there is a cat in front of him and that he has not got the slightest reason to think that appearance is deceptive here. Lacking any evidence to the contrary, he claims to be entitled to trust appearance. But now John replies: “You are right. You are entitled to trust appearance – but only in privileged conditions. So do you know that conditions are privileged?” Of course Tom could counter by pointing out that in lack of any countervailing evidence, he is entitled to take conditions to be privileged. But that is unsatisfactory for obvious reasons. Consequently, appeal to defeasible generalizations seems to be of little help with the skeptic. We would only embark on an infinite regress. So what is needed instead is an entitlement, which is not subject to further qualification (– and which can therefore help to stop the regress). Of course an entitlement can be lost. But that does not show that one has to do something to earn it in the first place. Nor does it show that the entitlement is admitted only if conditions are privileged.

And there is another reason why defeasible generalizations are not sufficient for justification. Suppose appearance is trustworthy in a particular case. The wall looks red to you. And it is indeed red. But suppose further that a well-informed friend told you that there is red light installed in the building. As it happens, the information is false. Nevertheless, you wouldn’t be justified in believing that the wall is red unless you could rule out that it is just illuminated so as to look red. So it is not enough that appearance is trustworthy in a given case. Being justified in a given belief is not just a question of whether the world is cooperating so as to make your belief come out true. Being justified requires being epistemically responsible. And epistemic responsibility is not reliability. It is a question of being able to properly respond to all the requirements of the situation and to see the force of undermining considerations.

But if it is not by appeal to defeasible generalizations, how else can we justify the claim that at least some of our presuppositions are legitimate? The short answer is this: To say that a presupposition is

legitimate in a given context is to say that a reasonable and informed observer could not but make it: given the purpose and direction of the inquiry in the context, it is rationally mandatory to make it. For lack of space I can only very briefly indicate the direction in which a more detailed answer might be found.

Usually, we presuppose a lot of things. Firstly, there are those propositions to doubt which would be to doubt the significance and intelligibility of the whole epistemic enterprise. We might call them, following Michael Williams, intelligibility or semantic constraints (Williams, 2001, p. 159). We don't have any evidence for them or any idea how such evidence would look like (cf. Wright, forthcoming). For a more detailed account of what is at issue here.) Secondly, we presuppose certain propositions whenever we raise a question or start an investigation. Let us call them, again following Michael Williams, methodological necessities (Williams, 2001, p. 160). They are those assumptions we cannot but take for granted if the investigation is to get off the ground at all. (Note, by the way, that to say that we presuppose something is not to say that we do anything. It is not that we deliberately adopt a certain attitude towards a particular set of propositions. Rather, we lack any such attitude.)

Now why should these presuppositions be legitimate? Very roughly, the idea, familiar from the writings of Wittgenstein, Wright and others, is that these presuppositions are legitimate in that they are necessary if we are to engage in any epistemic enterprise at all. As Wittgenstein puts it in §341: "... the questions that we raise and our doubt depend on the fact that some propositions are exempt from doubt, are as it were like hinges on which those turn" (Wittgenstein 1969). (A somewhat similar idea is Mark Timmons' idea of contextually basic beliefs, cf. Timmons, 1996.)

So we are entitled to take certain things on trust because otherwise we had to abandon all our epistemic enterprises anyway. We could not even begin to ask questions. So in these cases we have a *default entitlement*. Of course, as I said before, the entitlement can be lost. (And of course a fuller account would have to list and explain the differences between the various kinds of presuppositions and the corresponding default entitlements.) And we keep the entitlement only if we are epistemically responsible. Also, we keep it only if we are not guilty of ignoring something, which, if brought to our attention, we would clearly, and reasonably have to consider relevant to the investigation at hand. (Think of Henry driving through barn-façade-country.) But this, again, does not show that we have to do something to earn the entitlement in the first place.

So let us take stock. Lance and Little take MP to be more or less analogous to EH. Accordingly, EH seems to inherit all the problems of MP or EP respectively. Now my point was that (i) the tie between EH and EP should be severed, so that (ii), even given EH, we can make room for epistemic principles which do not take the form of defeasible generalizations. And if that is, at least roughly, on the right track, then there seems to be no need for defeasible generalizations in epistemology. They don't seem to add any explanatory power to our theory or to play any role in justification. This is not to deny that they say something true or might prove very useful elsewhere. It is just to deny them any explanatory or justificatory relevance in epistemology.

ACKNOWLEDGMENTS

I want to thank the organizers and the participants of the conference *Contextualist Approaches to Epistemology* in Mainz, September 2003 for making it such an interesting and stimulating conference. I especially want to thank Mark Lance, Christoph Jäger, Oliver Scholz, and an anonymous referee for this journal for very helpful comments on an earlier version of the paper.

REFERENCES

- Dancy, J.: 1993, *Moral Reasons*, Blackwell Publishers, Oxford.
- Dancy, J.: 2001, 'Moral Particularism', in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, (Summer 2001 Edition), URL = <http://plato.stanford.edu/archives/sum2001/entries/moral-particularism/>.
- Hooker, B.: 2000, 'Moral Particularism: Wrong and Bad', in B. Hooker and M. Little (eds.), *Moral Particularism*, Clarendon Press, Oxford, 1–22.
- Lance, M. and Little M.: 2004, 'Defeasibility and the Normative Grasp of Context', *Erkenntnis* **61**, 435–455.
- McDowell, J.: 1997, 'Virtue and Reason', in R. Crisp and M. Slote (eds.), *Virtue Ethics*, Oxford University Press, Oxford, 141–162.
- Pollock, J.: 1986, *Contemporary Theories of Knowledge*, Rowman and Littlefield, Totowa NJ.
- Stalnaker, R.: 1998, 'On the Representation of Context', *Journal of Logic, Language, and Information* **7**; reprinted in Stalnaker, R.: 1999, *Context and Content*, Oxford University Press, New York and Oxford, 96–113.
- Timmons, M.: 1996, 'Outline of a Contextualist Moral Epistemology', in W. Sinnott-Armstrong and M. Timmons (eds.), *Moral Knowledge*, Oxford University Press, Oxford and New York, 293–325.

- Williams, M.: 2001, *Problems of Knowledge: A Critical Introduction to Epistemology*, Oxford University Press, Oxford and New York.
- Wittgenstein, L.: 1969, *On Certainty*, Basil Blackwell, Oxford.
- Wright, C.: Forthcoming, 'Wittgensteinian Certainties', in D. McManus, (ed.), *Wittgenstein and Skepticism*, Routledge, London and New York.

Westfälische Wilhelms-Universität Münster
 Philosophisches Seminar
 Domplatz 23
 48143 Münster
 Germany
 E-mail: nkompa@uni-muenster.de

HANS ROTT

STABILITY, STRENGTH AND SENSITIVITY:
CONVERTING BELIEF INTO KNOWLEDGE

ABSTRACT. In this paper I discuss the relation between various properties that have been regarded as important for determining whether or not a belief constitutes a piece of knowledge: its stability, strength and sensitivity to truth, as well as the strength of the epistemic position in which the subject is with respect to this belief. Attempts to explicate the relevant concepts more formally with the help of systems of spheres of possible worlds (à la Lewis and Grove) must take care to keep apart the very different roles that systems of spheres can play. Nozick's sensitivity account turns out to be closer to the stability analysis of knowledge (versions of which I identify in Plato, Descartes, Klein and Lehrer) than one might have suspected.

1. INTRODUCTION: GRADES OF KNOWLEDGE AND BELIEF

Gettier has shattered our understanding of knowledge. There is still little agreement among philosophers what knowledge is. Stability theories (also known as defeasibility theories) say that knowledge is belief with a stable (indefeasible) justification. Nozick advanced an influential theory, according to which knowledge is belief that is sensitive to truth (or that "tracks truth"). In the contextualist model of Keith DeRose knowledge depends on how strong the subject's epistemic position is with respect to the belief in question.

It is well known that these proposals have difficulties in dealing with certain classes of counterexamples.¹ But my aim in this paper is not to confront the various theories with yet more and yet more complicated examples and counterexamples. I rather take it that they all capture important intuitions that can in some way or other be regarded as relevant to the question whether or not a given belief constitutes a piece of knowledge. The questions I am going to address are the following: Can stability, or more exactly, the stability of beliefs in an interrogation with a truthful critic like Socrates in a Platonic dialogue, be the right basis for the explication of knowledge? Does strength of belief imply stability, or *vice versa*? If knowledge lies in the stability of a belief, how does knowledge relate to the strength



of the belief? How does the strength of the subject's epistemic position with respect to a belief as highlighted in the contextualist literature relate to the strength of the subject's belief? And finally, as the contextualist account is at least in part inspired by Nozick's truth-tracking or sensitivity account of knowledge: How does the sensitivity of a belief for truth relate to the strength of the belief?

In my attempt to answer these questions I shall make use of a possible worlds modelling for subjunctive conditionals going back to Lewis (1973) and referred to by Nozick and DeRose. I shall draw attention to the fact that the same formal model can be used for the analysis of the stability approach, but that the interpretation of this model must then be crucially different. It represents the subject's doxastic state and can be used to represent the changes that this state undergoes while a critic tries to undermine the subject's beliefs by advancing potential defeaters. The possible worlds model can thus be used to model both internal (subjective) and external (objective) aspects of knowledge. I close with a short overview of the relation between stability, strength, epistemic position and sensitivity.

2. THE STABLE BELIEF THEORY OF KNOWLEDGE

The first stability account of knowledge is probably to be found in Plato's *Meno*, where Socrates says that true beliefs convert into knowledge if and only if they become "permanent" after having been "tied down" by giving reasons for them.² A less widely known but similar formulation is given by Descartes in his second Replies who claims that true knowledge cannot be "rendered doubtful".³

After Gettier's seminal 1963 paper, the idea of stability or indefeasibility has always loomed large in epistemological discussions. For instance, Klein (1971, p. 61) suggested the following *Felicitous-coincidence principle*:

If *S*'s evidence for α and a description of some of the particular circumstances in which *S* believes that α are such that it would not be reasonable to expect that α is true (based upon *S*'s evidence), even if α is true, *S* does not know α . Consequently, we might tentatively assert that *S*'s evidence for his belief that α is not sufficiently strong to certify his belief as knowledge if there is some fact which, were *S* to become aware of it, ought to cause *S* to retract his knowledge claim.

In the cases described by Klein, the critic just needs to point out to *S* the circumstances that make *S*'s belief that α unreasonable. This

should be sufficient to talk *S* out of believing α . Thus α is not a piece of knowledge according to the stability account.

From Lehrer (1965) at least up to Lehrer (1990), Keith Lehrer has been one of the most prominent champions of the stability account of knowledge. I shall concentrate on the version presented in Lehrer (1990, chapters 6 and 7). Like Plato, Lehrer suggests a dialogical construal of the stability idea. The believing subject is imagined as being engaged in a dialogue with a *critic*⁴ (a Socratic dialogue partner) who tries to undermine the subject's beliefs. Only if the subject *wins* the dialogue in the sense that he successfully defends his belief against all the critic's objections, can that belief be called *knowledge*.

Two of the essential rules of the *justification game* are that the critic is omniscient and that she confronts the subject only with information that is true. Such a test for knowledge may appear as a purely internal affair, since it seems to involve only the subject's beliefs and changes of belief, that is, only his internal states. But this is not quite true. The assumption that the critic's objections make use only of *true* statements brings in a connection with the actual world. Truth is what binds subjective beliefs to objective facts.⁵

So beliefs that fall short of knowledge are vulnerable. A point highlighted in many reactions to Gettier's examples is that the *justification for a belief* may be lost if new evidence comes in. Plato's original point, in contrast, was that *the belief itself* may be lost. This is a simpler idea, since it does not depend on the notoriously controversial concept of justification. Let us suppose that the belief changes occasioned by the incoming evidence are rational in some sense. Then, it seems, beliefs that persist enjoy some sort of justification. I want to make this simplifying assumption and base my discussion upon the following explication of knowledge:

A belief α is a piece of knowledge of the subject *S* iff α is not given up by *S* on the basis of *any true* information that *S* may receive.

This is what I will call the *stable belief theory* or the, shorter, *stability theory* of knowledge. My avoiding the term *defeasibility theory*⁶ is intended to mark terminologically the difference between the loss-of-justification and the loss-of-belief ideas. None of the approaches I am dealing with is based on the idea of justification. We presuppose that the subject is in some sense rational in accommodating his beliefs to new information, but we do not assume that justification plays a major role in such processes of belief adaptation.

3. NOZICK'S SENSITIVITY THEORY AS ENTAILING STABILITY

Nozick's (1981) influential *truth-tracking account* or *sensitivity account* of knowledge is usually presented as an important alternative to indefeasibility theories. But it is worth emphasizing that this account was devised so as to entail an element of stability as well. According to Nozick, a subject S knows that α if and only if (1) α is true, (2) S believes that α , and the following subjunctive conditionals are true:

$$(3) \quad \neg\alpha \Box \rightarrow \neg(S \text{ believes that } \alpha)$$

$$(4) \quad \alpha \Box \rightarrow S \text{ believes that } \alpha$$

Nozick's third condition is a *variation condition*, while (4) is an *adherence condition* (Nozick, p. 211). Regarding (3), the question to be answered is this: What would happen if α were false? (– which in fact it is not). The question regarding (4) is a little harder to formulate. Try this: What would happen if α were true? (– which in fact it is). It sounds strange to call (4) a *subjunctive* conditional even though α is known to be true. Like condition (3), condition (4) is supposed to have some modal force: “Not only is α true and S believes it, but if it were true he would believe it. . . . The truth of antecedent and consequent is not alone sufficient for the truth of a subjunctive” (Nozick, 1981, p. 176, variable renamed).⁷ What is particularly interesting for our topic is that the antecedent of condition (4) is supposed to cover α -worlds in which the subject is interrogated by a critic. Nozick himself relates (4) to the situation of the Socratic dialogues:

Meno claimed he could speak eloquently about virtue until Socrates, torpedolike, began to question him. He did not know what virtue was, for Socrates' questions uncovered Meno's previously existing confusions. Even if it had been a sophist's questions that bewildered Meno, getting him to believe the opposite, what he previously had would not have been knowledge. Knowledge should be made of sterner stuff.

Thus, some skeptical arguments play off condition 3, others off condition 4.⁸

It is clear that for Nozick the contingent truth of α and β in the actual world w_a does not suffice to make the subjunctive conditional $\alpha \Box \rightarrow \beta$ acceptable. But how far must we be ready to deviate from the actual course of events in order to test for the truth of the conditional? How far does the truth of β have to extend among the α -worlds?

In his attempt to answer this question, Nozick employed a model using *spheres of possible worlds* due to Lewis (1973). A sphere is the set of possible worlds that are similar to the actual world w_a up to a certain degree. The smallest sphere is the singleton $\{w_a\}$. We already said that this set is not enough for the evaluation of the conditional, we have to consider larger spheres. But is it sufficient to consider the second smallest sphere, the set of possible worlds that are closest to, but not identical with the actual world; or do we have to go out until we meet the closest $\neg\alpha$ -worlds; or is some intermediate sphere adequate? Nozick (1981, pp. 680–681) has a long, complicated and somewhat irresolute footnote about this question, suggesting that we must indeed go out to a level that includes at least the closest (but maybe many more) $\neg\alpha$ -worlds.⁹ Taken together, it seems that (3) and (4) are meant to imply that the subject's belief-that- α covaries with the fact-that- α “for some distance out in the closest α band to the actual world”. Since Nozick thinks that this band contains worlds in which critical conversations with the critic take place, we conclude that meeting a critic does not mean a big deviation from the actual world for Nozick.¹⁰

So there is a lot of support for the stability analysis in the epistemological literature. Two things remain to be noted. The approach does not take care of the case where the subject is presented with *misinformation*. It is not clear whether knowledge should be robust against local errors of perception or memory, or wrong testimony from the critic. To take up Nozick's phrase, shouldn't knowledge be made of still sterner stuff – stuff that also survives (a modest amount of) misinformation? I just want to raise the question here; I am not going to further pursue it in this paper.

Secondly, even true information may be misleading. Sometimes there is a definite bias in the kind of information that we receive (from a used-cars salesman, for instance). Even if every single piece of information the subject receives is true, the picture that emerges may tempt him to draw the wrong inferences, thereby undermining what he (apparently) knew before. The problem of misleading defeaters and pseudo-defeaters of knowledge has accompanied stability theories from their beginning, and we will return to this point below.

4. AN INTERNAL AFFAIR: STRONG BELIEFS

One may plausibly expect that the stability of a belief derives from its strength. It is instructive to look at the relation between these

concepts more closely. We have to account for varying degrees of belief, and we will do that in the simplest possible way, by means of a qualitative modelling.¹¹ Let us look at two interdependent ideas to represent the idea of *strong belief*:

- (a) high epistemic entrenchment (high epistemic rank);
- (b) stability (persistence, tenacity) in certain kinds of belief change.

As a model for belief states we take a subjectivist version of the model already appealed to by Nozick. Formally, we replace Lewis's (1973) objectivist conception by Grove's (1988) subjectivist conception of systems of spheres.¹² Let us represent a doxastic state by a system of nested sets of possible worlds, supposing, for the sake of simplicity, that everything is finite. The smallest set is the set of possible worlds which the subject believes to contain the actual world w_a . If the subject receives evidence that the actual world is not contained in this smallest set, he falls back on the next larger superset. And again, should it turn out that the actual world is not to be found in this set either, the subject is prepared to fall back on the next larger set of possible worlds. And so on. The sets or *spheres* of possible worlds correspond to spheres of plausibility, or to put it differently, spheres of deviation from the subject's beliefs. The spheres are the subject's personal spheres of possible worlds as it were, spheres for the first person. The system of spheres taken as a whole represents a mental state (viz., a doxastic state) and must not be expected to be centered on a single world w_a that represents the actual world. If one of the subject's beliefs is wrong, then w_a is not even contained in the innermost sphere, but may occur at any arbitrary position in the sphere system. Let us now see how we can use this modelling to represent the idea of strong belief.

Re (a) We can identify the strength of a belief with its degree of *doxastic entrenchment*, where the degree of doxastic entrenchment of a belief α can be measured by the number of spheres that contain exclusively α -worlds. The more spheres (i.e., the more fallback positions) are fully covered by α , the better entrenched α is.¹³

Re (b) The entrenchment terminology suggests that we are interested in how hard it is to eliminate a belief. Rather than defining the resistance against elimination with reference to a fixed doxastic state, we can refer directly to the potential developments of that doxastic state. A belief is *stable* to the extent that it is unlikely that the belief is lost in processes of belief change.

What kinds of belief change should we take into account? Here we return to the dialogue model with the critic, and add that a third rule

of a Lehrerian justification game is this: The subject must accept the pieces of true information the critic provides it with (in this sense, the critic's objections must be "successful"). So the subject has to be ready to actually *add* new information. Two questions suggest themselves: Should we be ready to account for the case where what the critic tells the subject is *incompatible* with the latter's beliefs? Should we be ready to account for the case where the critic prompts the subject to *subtract* a belief rather than add a new one? It is important, I am going to argue now, that both questions are answered in the affirmative.

The need for *belief-contravening revisions*, belief changes induced by new information that contradicts the subject's prior beliefs, is obvious if we endorse a simple thesis of *fallibilism*: For all subjects and at all times, some of the subject's beliefs are wrong.¹⁴ That we are all fallible is a basic fact of life. Human beings have a hard time refraining from believing, they tend to be credulous, and many people think: excessively credulous. As a consequence, we always have to face the fact that some of the countless beliefs we hold are mistaken. As we gather more evidence and obtain more true information from various sources (from our relentless critic, for example), we will sooner or later encounter conflicts with our previous beliefs. In such cases, we have to perform belief-contravening revisions.

The story of how to base a revision of the subject's beliefs upon the system of spheres representing his doxastic state is easy to tell (Grove 1988, Gärdenfors 1988). If β is the new bit of information, the subject looks for the smallest sphere that contains at least one β -world. The subject believes α after successful performance of a revision by β just in case α is true in all β -worlds that are contained in this smallest β -admitting sphere. This recipe works regardless whether β is or is not consistent with the subject's previous beliefs. Figure 1 may serve as an illustration.

5. THE STABILITY ANALYSIS OF KNOWLEDGE AND THE STRENGTH OF BELIEFS

One will be inclined to think that there must be a tight connection between the strength of a (true) belief and its stability under (truthful) criticism. But the two concepts cannot be identical since strength seems to be a purely "internal" property, whereas stability as just defined imposes "external" constraints through the requirement that the critic's statements be all true.




Figure 1. Here, α is a true belief, but α (or more exactly, $\beta \supset \alpha$) is not sufficiently well-entrenched to survive the revision by the true information β . So S does not know that α .

The stability account may be formulated in the setting of one-shot belief revision: Subject S *knows that* α if and only if S believes that α and α is not given up by S after receipt of any true information (from the critic, say). More precisely, α is a piece of knowledge of S if and only if α is not lost when S 's set of beliefs is revised by any arbitrary true piece of information. With a little help from belief revision theory, we will now prove the following result:

OBSERVATION. *The belief α is stable with respect to the revision of S 's belief set by any true piece of information if and only if α is more entrenched in S 's belief state than every false belief; or equivalently, in the system of spheres modelling: if and only if α holds not only throughout the innermost sphere but also throughout the smallest sphere containing the actual world w_a .*

Proof. We want to show that α is stable under truthful revision iff it is more entrenched than any falsehood, in symbols:

$$(\dagger) \quad \forall \beta (\beta \text{ is true} \Rightarrow \alpha \text{ is in } B^* \beta) \quad \text{iff} \quad \forall \gamma (\gamma \text{ is false} \Rightarrow \gamma < \alpha)$$

where B denotes the subject's original belief set and $B^* \beta$ denotes the belief set that results from revising B by the sentence β .

First of all we have to connect the notion of entrenchment with the subject's belief change behaviour. A sentence β is at most as *entrenched* as a sentence γ , in symbols $\beta \leq \gamma$, iff β is lost when the subject learns that the conjunction of β and γ is not true, in symbols, iff β is not in $B^* \neg(\beta \& \gamma)$. Call this the definition of entrenchment.¹⁵

This definition entails the dominance condition which says that $\beta \leq \gamma$ whenever β logically implies γ (in this case β is not in $B^* \neg(\beta \& \gamma) = B^* \neg \beta$).

The left-hand side of $(\dagger)$ implies the right-hand side: Suppose the right-hand side is false, i.e., α is not more entrenched than every

falsehood. Then there is a false γ such that $\alpha \leq \gamma$. Now consider $\neg(\alpha \& \gamma)$. This sentence is true, since γ is false. By $\alpha \leq \gamma$ and the definition of entrenchment, it follows that α is not in $B^*\neg(\alpha \& \gamma)$, so α is not stable under truthful revision, i.e., the left-hand side is false.

The right-hand side of $(\dagger)$ implies the left-hand side: Suppose the left-hand side is false, i.e., α is not stable under truthful revision. Then there is a true β such that α is not in $B^*\beta$. Since β is in $B^*\beta$ and this set is logically closed, it follows that $\beta \supset \alpha$ is not in $B^*\beta = B^*\neg((\beta \supset \alpha) \& \neg\beta)$ (notice that $\neg((\beta \supset \alpha) \& \neg\beta)$ is logically equivalent with β). By the definition of entrenchment, this means that $\beta \supset \alpha \leq \neg\beta$. Now by the dominance condition, $\alpha \leq \beta \supset \alpha$, so by the transitivity of entrenchment $\alpha \leq \neg\beta$. Since $\neg\beta$ is false, we have found a falsehood that is at least as entrenched as α , i.e., the right-hand side is false. $\square$

Using this Observation, we can see that knowledge in the stability interpretation does not require maximal entrenchment,¹⁶ but it is indeed characterized by a certain degree of entrenchment (i.e., by a certain strength of belief). The particular strength of belief that is required depends on the position of the actual world w_a in the system of spheres. If the subject considers w_a to be a fairly plausible world, knowledge does not require very strong belief. If, however, w_a is far out in the subject's system of spheres, knowledge requires very highly entrenched belief. *Prima facie*, it looks like an element of epistemic luck where in the subject's system of spheres the actual world happens to be placed. But perhaps it is not luck after all where w_a is being located, but rather merit – a sign of how good S 's doxastic state is. It is certainly a virtue of an epistemic subject to have his beliefs in good accord with the actual world.

Still I think that the Observation discloses a problematic feature of the stability analysis which ties knowledge too tightly to the strength of belief. As a first indication, consider the epistemically ideal case in which S 's beliefs are all true. In the sphere model, this means that w_a is contained in the innermost sphere. Then, according to our Observation, a truthful critic can never talk S out of believing *any* of his beliefs. True information will only result in a consistent addition of beliefs (i.e., in the elimination of possible worlds from the innermost sphere). According to this analysis, if all of the subject's beliefs are true, each and every belief of his constitutes a piece of knowledge. This, however, is counterintuitive. Intuitively, having only true beliefs does not protect S against being dissuaded from believing a particular one of his beliefs. Problems for the more realistic case where S has some false beliefs will be discussed in Section 8.

6. CRITICS, SKEPTICS AND THE MEANING OF *MIGHT* SENTENCES

The skeptic is not so much a provider of new evidence as someone who *raises doubts* and *calls beliefs into question*. The critic, we said, supplies the subject with new, truthful information. The skeptic, in contrast, does not furnish positive information. Her mission is a negative one, it typically leads to the subject's relinquishing some information without getting anything new. That is, she instigates processes of belief *elimination* or *contractions* of belief sets rather than their revisions. The skeptic does not positively claim that *S* is a brain in a vat, she rather points out that *S might be*, for all he knows, a brain in a vat. She does not assert that those animals in the zoo of Berlin *are* cleverly painted mules, she only says it is *possible* that they are.

This leaves us with the question of how to deal with such modalized statements. Assuming again that the subject *S* has to accept what the skeptic is saying, we need to specify the sort of belief change that goes on in *S* after accepting the skeptic's *might* sentence. So suppose the skeptic says *might- α* . What the subject does first, I suggest, is try out what his beliefs would look like after accepting α . But then, since he has no positive evidence that α is actually true, he settles for what is common to his current belief set *B* and the result of revising *B* by α . This procedure can also be reinterpreted as a process of withdrawing $\neg\alpha$ from the subject's belief set *B*; in this reading "revise *B* by *might- α* " means "withdraw $\neg\alpha$ from *B*".¹⁷

If we admit *might* sentences as skeptical objections, we can frame an argument to the effect that Nozick's positive conditional (4) implies his negative conditional (3), given that α is true. We show this by contraposition. So suppose that α is true, but that *not*

$$(3) \quad \neg\alpha \Box \rightarrow \neg(S \text{ believes that } \alpha)$$

According to the semantics for subjunctive conditionals, this means that there is a relevant possible world such that

$$(\ddagger) \quad \neg\alpha \text{ and } S \text{ believes that } \alpha$$

is true in that world. Now assume that the critic tells the subject about this possibility by uttering the sentence

$$\textit{might}-(\neg\alpha \text{ and } S \text{ believes that } \alpha)$$

According to the rules of the justification game, *S* accepts this sentence. We said that this means that *S* checks, for the sake of argument, what his beliefs would look like after revising them by $(\ddagger)$. Since consistency is to be respected, the subject loses his prior belief α

in the revised belief set, and α remains lost of course if this set is intersected with the original belief set. We have now described a scenario in which α is true and “ S believes that α ” is false. Let us assume (with Nozick) that such a scenario is plausible and relevant, and thus close to the actual world. Then it follows that the positive conditional

$$(4) \quad \alpha \Box \rightarrow S \text{ believes that } \alpha$$

does not hold. This completes the proof that (4) implies (3), provided that α is true, the conversation with the skeptic is close to the actual world and the skeptic is allowed to put forward *might* sentences.

What difference does it make whether the “information” supplied by the critic comes in the form of a categorical or in the form of a modalized sentence? It makes a big difference, since the rules of the game constrain her to give true information only. If the sentence α is false, then *might- α* may still be true. So the critic – or rather: the skeptic – has a lot more possibilities to talk S out of believing a proposition if she is allowed to use *might* sentences. By assumption, she is omniscient, she knows the whole truth, and she speaks nothing but the truth. But when is a *might* sentence true? This, of course, depends on the meaning of the modality. The most common reading of skeptical objections is to understand *might* epistemically: S cannot exclude, for all he knows, that he is deceived by an evil demon, that he has been envatted by evil scientists, that this animal in the zoo is a cleverly painted mule etc. But it is doubtful that the epistemic understanding of *might* is the right one to plug in into Nozick’s conditions (3) and (4). We are stepping into deep waters, waters that we cannot even begin to fathom out here. For the rest of this paper, I shall assume that the objections that the critic raises can be expressed in non-modal terms. I will not deal with *might* sentences any more. My critic is not supposed to be a skeptic.

7. MORE ON INTERNAL AFFAIRS: DIALOGUES AND PIECEMEAL EVIDENCE

Plato’s Socrates liked to stretch his teaching out in long dialogues. Lehrer, too, used dialogues to illustrate his concept of knowledge. It has been argued that there is not only a heuristic, but an epistemologically significant difference between presenting corrective evidence all at once and presenting it *seriatim*.¹⁸ In order to account for this, the classical model of one-shot belief revision must be extended to a more elaborate one. A conversation with the critic typically consists

of several rounds, in each of which she would release new information. A good model of the belief change that the subject is experiencing in such a conversation must be able to describe *iterated belief changes*.

If we want to stick to the simple systems of spheres modelling, there is a rather limited number of methods for iterated belief change, and is not quite clear which (if any) of these models can adequately capture the kind of process that we need for the conversation with the critic. Consider for illustration a slightly modified variant of an example of Lehrer (1965). Let p stand for the sentence “Jones owns a Ferrari”, q , r , s and t for corresponding sentences about other colleagues of the subject owning a Ferrari. Let us suppose that the doxastic state of Gettier’s subject regarding this matter is represented by the system of spheres in Figure 2.

The subject’s initial beliefs include that Jones has got a Ferrari, while the others have not. S thinks that w_1 is the actual world. The second most plausible situation is the one in which none of his friends owns a Ferrari, i.e., w_0 . Only at the next level are there worlds in which some of his other colleagues owns a Ferrari: worlds w_2 through w_5 . Figure 2 does not show the still more far-fetched situations in which more than one of his friends owns a Ferrari. After all, Ferraris are not meant to be everybody’s cars.

Now assume that it is in fact Brown who owns a Ferrari ($w_a = w_5$, say), and imagine the critic beginning to tell S the truth about the situation. Her first hint is

- (1) “Jones has not got a Ferrari.” ($\neg p$)

The subject’s straightforward reaction is, on any of the standard accounts of belief revision, to proceed to a belief state that takes w_0 to




Figure 2. Gettier case, with S believing that Jones owns a Ferrari (p).

be the true world. With this, S is still wrong. Imagine the critic passing on a second piece of information to the subject

- (2) “Someone has got a Ferrari.” $(p \vee q \vee r \vee s \vee t)$

Now the classical one-shot belief revision theory of the 1980s (Gärdenfors 1988) was at a loss about how to revise the subject’s beliefs in the second step. In the 1990s, however, a number of techniques were developed to deal with iterated changes in the simple possible worlds setting that we are using in this paper.¹⁹ Different methods lead to different reactions to the second of the critic’s hints. If S chooses to apply the method of *conservative belief revision*, he returns to w_1 as the most plausible world, and again believes that Jones owns the Ferrari. However, since our critic invariably tells the truth, forgetting about (1) is not the type of reaction that we would like to see. If S applies the method of *moderate belief revision*, he reaches the conclusion that the true world is among $w_2, \dots, w_5$, and thus believes that the owner of the Ferrari is one of the persons in question, with the exception of Jones. This conclusion is what we expect of a rational person.

If knowledge is stable belief, Jones cannot be said to know that someone in his class owns a Ferrari – which is in accordance with our intuitions. If the revision method employed is the conservative one, however, then Jones may be said to know that *if* someone owns a Ferrari, then it is Jones. This is too conservative. The subject should be able to learn more from the critic’s information, he should not revive in the second step his false initial belief that Jones is the owner of the Ferrari. The method of moderate belief change is just what we need for the dialogue with the critic. It consistently accords incoming information priority over old beliefs. In fact, since everything the critic says is true (by the rules of the justification game), the conjunction of her statements is consistent. Moderate belief change is such that iterated changes by a sequence of jointly consistent bits of information $\alpha_1, \alpha_2, \dots, \alpha_n$ always result in the same belief set as a single change effected by the conjunction $\alpha_1 \& \alpha_2 \& \dots \& \alpha_n$.²⁰ All evidence supplied by the critic *seriatim* can be collected and has the same effect as if the evidence were presented all at once.

We can conclude that belief revision theory has the resources appropriate to deal with a stepwise correction of the subject through an extended dialogue with the critic. As long as each piece of input is true, the stability of a belief under sequences of revisions is reducible to its stability under various one-shot revisions. The Observation of Section 5 linking the stability account to strengths of belief transfers to the iterated case without modification.

8. A PROBLEM FOR THE STABILITY ACCOUNT OF KNOWLEDGE

We have mentioned in the introduction that defeasibility theories of knowledge were diagnosed as problematic soon after their invention. We shall now show that our move to the stability theory (that substitutes loss-of-belief for loss-of-justification) does not get round the problems. The point is that it is fairly easy for the critic to talk the subject out of a belief, even if intuitively the belief constitutes genuine knowledge. Consider the following abstract argument due to Jacob Rosenthal which can actually be seen as an illustration of our Observation in Section 5.²¹ Suppose that S knows that α , but that α is not maximally entrenched in S 's belief state. Suppose further that S has a very well-entrenched belief β that happens to be false. Let us assume that α is not more entrenched than β . Then S can be talked out of believing α in the following way. The critic correctly points out that $\alpha \ \& \ \beta$ is false. By the rules of the justification game, S recognizes that what the critic says is right, and he accepts $\neg(\alpha \ \& \ \beta)$. In order to maintain the consistency of his beliefs, S has to remove $\alpha \ \& \ \beta$. Being logically competent, S realizes that he has to remove either α or β . By our hypothesis that α is not more entrenched than β , the belief α has to go (this is what the term "entrenchment" means). So S has been talked out of believing α by the critic. – Hence, if the stability analysis of knowledge is correct, S has not known that α to begin with. Contradiction. Hence S can know α only if α is more entrenched in S 's belief state than every other belief that happens to be false. One well-entrenched false belief erases as it were a lot of putative knowledge that has not got anything to do with it.

Now the obvious question is: Doesn't this show that the stability analysis is fundamentally flawed? Tentative answer: No, but we have to refine it. Intuitively, it seems the critic should only question statements that are somehow "basic", statements on which α depends rather than statements that depend on α themselves. And in the argument just sketched, the criticized proposition $\alpha \ \& \ \beta$ was presented as parasitic on the (more) basic beliefs α and β .

But we cannot get rid of the problem that easily, as is shown by the following more concrete example. Suppose I think I observed that Grabit stole a book from the library at 3 p.m. Suppose further that I had forgotten my glasses that afternoon. So, being short-sighted, I am not absolutely sure that it was Grabit who stole the book (p), although for all practical purposes I would not hesitate to rule out the possibility that it was someone else. When making this observation, I

looked at my very reliable Rolex watch, so I am very sure that it was 3 p.m. when the book was stolen (q). I have an excellent reason to believe q , a better reason anyway than I have for my believing that p is true. As a matter of fact, however, Grabit did steal the book, but it was already 3:30 when that happened. (My reliable Rolex had stopped working for a while, a fact that escaped my attention because it later reset itself with the help of a radio signal.) By everyday standards, I may truly be said to know that Grabit stole the book. But of course I cannot be ascribed knowledge that this event took place at 3 p.m. At that time Grabit was still having lunch with some of his colleagues, all respectable people who make for irreproachable witnesses. Now a critic may rightfully point out to me that my original belief that Grabit stole the book at 3 p.m. is not true. Being forced to retract this belief, I conclude, on the basis of the quality of the evidence that I possess, that p must be false and q must be true.

This is certainly a rational reaction. The critic, however, has managed to talk me out of believing something that I seem to have known before, viz., that Grabit stole the book (p). What are we to say now? Was p unstable knowledge, or was it no knowledge at all? The stability theorist is committed to saying it wasn't knowledge to begin with, but this seems counterintuitive. The mere fact that the subject has a false belief q that is sufficiently well-entrenched to drive out the true belief p should not in itself be sufficient to discredit p 's claim to the status of knowledge. But the tentative answer to the abstract case described before does not seem to be available any more. There is no reason to deny that my "basic" belief was precisely that Grabit stole the book at 3 p.m. Doesn't it look artificial to formalize this original belief by p & q ? There is no good motivation for splitting this belief up, as we just did for the sake of exposition, into the two halves "Someone stole the book at 3 p.m." and "Grabit stole the book some time". And it can hardly be claimed that the critic's clue was misleading.

As a side remark, the following observation may be interesting: It can be shown that if we assume that the subject's belief set is logically closed,²² then *no* belief-contravening revision that does not result in an omniscient belief set strictly enlarges the set of true beliefs. The subject is bound to lose some true information in his conversation with the critic, notwithstanding the fact that the latter's clues may improve the subject's belief set in any intuitive sense. Even if the subject's prior beliefs contain some falsehoods while the new piece of information as well as all posterior beliefs are true, the revision prompted by the critic will make the subject lose some true beliefs.²³ To be sure, this is only an observation about *belief*. But we begin to

form an idea that it is much easier for the critic talk the subject out of his *knowledge* than we might have suspected.

It still seems to me that the notion of stability captures an important aspect of knowledge, but I confess that I do not know how to repair the stability account so as to avoid the problems we have identified. We now leave the criteria that are internal in the sense that they refer to the development subject's mental state (without referring to processes of justification), and turn to the external notion of the strength of the subject's epistemic position.

9. EXTERNAL AFFAIRS

In the last two sections, I have used systems of spheres as representations of belief states, as structures that determine the strength of a subject's belief and help him to revise his beliefs not only once, but several times. Systems of spheres of possible worlds were also appealed to by epistemologists like Nozick (1981) and DeRose (1995). It is important, however, to keep distinct in the formal modelling what is distinct in substance. In Nozick's sensitivity account of knowledge, subjunctive conditionals are evaluated with the help of systems of spheres. These spheres are not those of the epistemic subject, but those of a third person that ascribes knowledge to the subject. They do not, however, represent the ascriber's belief state. They are meant to represent objective similarities between possible worlds. The location of a possible world in such a system does not represent the world's plausibility, but its distance (however conceived) from the actual world w_a . Nozick and DeRose appeal to Lewis-style systems centered on the actual world w_a . In contrast, the Grove-style systems used for the stability analysis are not centered on any world, and the position of the actual world w_a cannot be determined on *a priori* grounds (we said before that its placement may just be a matter of luck, but it may also be a sign for the aptness of the subject's doxastic state). Systems of spheres representing the subject's belief state are obviously subjective. Systems of spheres representing similarities are subjective in a much less evident way; similarity is always similarity in certain interesting or salient respects.²⁴

Systems of spheres may thus play very different roles in the analysis of belief and knowledge. There are also two different notions of *strength* that must not be confused. We have already noted that the *strength of a belief* may be identified with its degree of epistemic entrenchment in the subject's doxastic state. Everything about this

concept is internal, and the strength of a belief may be assumed to be completely transparent to the subject.

This is all very different from the *strength of the epistemic position* in which a subject is with respect to a certain belief. An interesting interpretation of this concept is suggested by DeRose (1995, pp. 490–492). A subject is in a *strong epistemic position* with respect to α if and only if his belief that α covaries with the truth of α not only in the actual world, but also in all worlds that deviate from the actual world to a quite significant degree. The more deviation from the actual world is tolerated without destroying the covariance of truth and belief, the stronger the epistemic position of the subject is. How strong or weak the subject's epistemic position actually is need not be transparent to him. The strength of the epistemic position is partly an objective matter (after all, it is the *actual world* that forms the center of the relevant system of spheres) and partly something to be judged by the third person's subjective standards (it is *her* similarity relation that serves as a measure of the deviation).

Might it be that a system of spheres is subjective and objective at the same time? Well, in cases of self-attributions of knowledge, the first person appears to take the role of the third person and many of the distinctions we have just made seem to collapse. But we should be aware of the fact that even when strength of belief (first person perspective) and strength of epistemic position (third person perspective) pull in the same direction, this is not sufficient to decide the case for knowledge. Let us have a look at the system of spheres depicted in Figure 3, which we now assume to represent both the subject's belief state and the attributor's similarity relation at the same time.²⁵ Let H stand for a skeptical hypothesis (like "I am a brain in a vat" or "This is a painted mule") and O stand for an ordinary hypothesis (like "I have hands" or "This is a zebra"). In the actual world, S believes that O , but does not believe that H is true. More importantly, since H and O are conceptually incompatible, the ordinary belief can only be true if the skeptical hypothesis is false.

In the situation depicted in Figure 3, S believes *more firmly* that $\neg H$ than that O , since $\neg H$ is true throughout four spheres and O is true only throughout two spheres (notice that the smallest sphere is the singleton set $\{w_a\}$). Similarly, the covariance of belief and truth extends farther with respect to $\neg H$ than with respect to O , which means that the S 's *epistemic position* with respect to $\neg H$ is *stronger* than with respect to O . Nevertheless, according to the sensitivity model of Nozick, the subject *knows* that O but does *not know* that $\neg H$. In each of the closest (most plausible) worlds where O is false, S would cease




Figure 3. Hypotheses and beliefs in hypotheses.

to believe that O , but there are some closest (most plausible) worlds where H is true and yet S would not believe it. Since the skeptical hypothesis is very far-fetched (as skeptical hypotheses typically are²⁶), the subject needs to be in a *very* strong epistemic position with respect to H in order to *know* whether it is true or false. Since on the other hand the ordinary hypothesis is much more mundane, the epistemic position with respect to O need not be very strong in order to know whether O is true or false. So even strength of belief and strength of epistemic position taken together provide no reliable indication of a belief's claim to the status of knowledge – if knowledge is understood as characterized by Nozickian truth tracking.

10. CONCLUSION

There are various routes to explicating knowledge without reference to the concept of justification. In this paper I have had a first look at the relationship between some other properties that have been thought to contribute to converting true belief into knowledge: stability, sensitivity to truth, strength of belief and strength of epistemic position. I have tried to make clear the different roles that can be played by a model using systems of spheres of possible worlds, and to sort out some subjective and objective factors involved.

While it is fairly obvious that strength of belief cannot in itself be a criterion for knowledge, the other properties have indeed been held to

be good criteria of knowledge, at least by some authors. It is time to summarize what we have found out about their mutual relationship.

The relation between sensitivity and strength of epistemic position has been set out nicely by DeRose (1995, pp. 491–492), and there is little to add to that. For ordinary beliefs, a *good* epistemic position is sufficient for knowledge according to the sensitivity analysis. For extraordinary beliefs (like the belief that a skeptical hypothesis is false) a good epistemic position is not normally sufficient; it has to be *excellent*. Conversely, if one knows that α according to the sensitivity analysis, this implies that the epistemic position is fairly good for ordinary beliefs, and it implies that it is excellent for extraordinary beliefs. Figure 3 shows how one's belief in O can be sensitive and one's belief in $\neg H$ can be insensitive while at the same time one's epistemic position is stronger with respect to $\neg H$ than it is with respect to O .

Now we turn to the stability of a belief which we saw to be intimately linked to its strength in Section 5. Our discussion has suggested that stability is related to sensitivity, but that this relation is far from perfect. Nozick himself pointed out that his fourth condition is meant to imply persistence of the belief under Socratic criticism. If the subject's belief cannot survive such a procedure, we have to deny the subjunctive conditional $\alpha \Box \rightarrow (S \text{ believes that } \alpha)$. Nozick's argument depends crucially on the idea that such a critical conversation may take place in worlds that are close to the actual world.

For the converse direction, it seems that knowledge according to the stability analysis implies knowledge according to the sensitivity analysis only in one of two possible cases. As we have seen in Section 6, if the negative conditional

$$\neg\alpha \Box \rightarrow \neg(S \text{ believes that } \alpha)$$

is wrong, then the critic can feed the subject with the modalized information that S *may* be in a situation in which S believes that α is true, even though $\neg\alpha$ is true. The subject's belief in α would then appear to be shaken. If, on the other hand, the positive conditional

$$\alpha \Box \rightarrow (S \text{ believes that } \alpha)$$

is wrong, then I cannot see how this could be detected by the critic's attempt to talk S out of believing α .

Regarding the relation between stability and strength of epistemic position (in the DeRose's sense), the latter implies the former in so far as α is true, given that we endorse Nozick's assumption that the

conversation with the critic takes place in the vicinity of the actual world. Stability, on the other hand, does not seem to imply strength of epistemic position.

The traditional notion of justification plays no role in the accounts that we have discussed. In fact, it is controversial even between the founders of the standard belief revision paradigm to what extent this paradigm can account for the justificatory structure of beliefs. Gärdenfors (1990) argues that foundationalist intuitions can be captured in the AGM model at least by reconstruction, while Makinson (1997) emphasizes how important it is to realize that this model does *not* come equipped with any justificatory structure. I tend to think that Makinson's picture better captures the nature of belief revision theory.

Belief revision theory thus seems orthogonal to the traditional concerns of mainstream epistemology. Perhaps the best account of how to understand their relation is still to be found in Harman (1986).²⁷ But there is also a fully developed alternative philosophical theory of knowledge that is not only congenial with belief revision theory, but has to some extent even motivated it, namely the work of Levi (1980, 2004). Levi's pragmatist attitude is strongly opposed to any kind of "pedigree epistemology" and importantly characterized by the thesis that it is not *beliefs* but *changes of belief* that are in need of justification, and that such justification has to be given in decision-theoretic terms. It seems to me, however, that Levi's account is still only loosely connected with mainstream epistemology. This is a regrettable state of affairs, and one that should be finished soon.

ACKNOWLEDGMENTS

This paper is a substantially revised version of a talk given at the Mainz conference on contextualism in September 2003. I want to thank Elke Brendel, Christoph Jäger, Michael Williams, and many other members of the audience for a stimulating discussion. Special thanks go to my commentator Lydia Mechtenberg and an anonymous referee of this journal for their valuable comments.

NOTES

¹ For critical discussions of stability theories, see Shope (1983, pp. 45–74) and of Nozick's theory, see the papers collected in Luper-Foy (1987).

- ² "True opinions too are a fine thing and altogether good in their effects so long as they stay with one, but they won't willingly stay long and instead run away from a person's soul, so they're not worth much until one ties them down by reasoning out the explanation. . . . And when they've been tied down, then for one thing they become items of knowledge [*ἐπιστήμη*], and for another, permanent [*μόνιμοι*]. And that's what makes knowledge more valuable than right opinion, and the way knowledge differs from right opinion is by being tied down." (*Meno*, 97e–98a; Plato, 1994, p. 69)
- ³ "But I maintain that the awareness [*cognitio*] of his [of the atheist, HR] is not true knowledge [*scientia*], since no act of awareness that can be rendered doubtful [*quae dubia reddi potest*] seems fit to be called knowledge." (Adam-Tannery edition, Vol. VII, p. 141; Descartes 1641/1984, p. 101).
- ⁴ Lehrer (1990) talks of a sceptic, who is renamed into the critic in Lehrer (2000).
- ⁵ I am neglecting here the constraint characteristic of Lehrer that the critic may only advance information about which the subject has had a definite belief to begin with. – In Rott (2003b), I have advocated a "dynamic" interpretation of the account presented in the first edition of Lehrer's *Theory of Knowledge* (1990). Lehrer (2003, p. 344) denies that his theory was ever meant to be dynamic. It seems to me, however, that Lehrer's (1990) continual talk of "moves", "rounds" and "combinations" of eliminations and replacements in the "ultra justification game" between a claimant and his critic clearly suggests an extended conversation with repeated turntaking. It is much harder to find anything dynamic in the substantially revised theory of the second edition of the *Theory of Knowledge* (Lehrer, 2000). Still, Spohn (2003) offers a sophisticated reconstruction of the new theory in terms of his belief change model.
- ⁶ Usually associated with people like Keith Lehrer and Thomas Paxson jr., Peter Klein, Marshall Swain, David Annis, Gilbert Harman and John Pollock.
- ⁷ Williams gives a good gloss of the conditional (4): "If, in somewhat changed circumstances, it were still the case that α , I should still believe that α ." (Williams, 2001, p. 30, variable renamed)
- ⁸ Nozick (1981, p. 213). This passage shows, I believe, that classifying Nozick as a pure externalist would miss an important point.
- ⁹ This is very similar to the *sphere of epistemically relevant worlds* as determined by DeRose's (1995, p. 493) *Rule of Sensitivity*: "When it's asserted that *S* knows (or does not know) that *P*, then, if necessary, enlarge the sphere of epistemically relevant worlds so that it at least includes the closest worlds in which *P* is false." (At this point I neglect that both Nozick and DeRose qualify their definitions by holding fixed a certain method of belief-acquisition; compare footnote 24 below.) Goldman (1987) gives arguments to the effect that the subject need not always go out that far for knowledge.
- ¹⁰ Nozick's assumption that meeting a critic or a skeptic is a nearby possibility is of course compatible with DeRose's presupposition that the truth of the skeptical hypothesis itself is a remote possibility.
- ¹¹ One can retain the spirit of the possible worlds modelling and in addition take advantage of the structure of ordinal numbers, thereby gaining a lot of additional expressive power. See Spohn (1988).
- ¹² More generally, one could use non-nested systems à la Lindström and Rabinowicz (1991).
- ¹³ Compare Lindström and Rabinowicz (1991) who incidentally also introduced the fallback terminology.

- ¹⁴ This broadly Peircean or Popperian notion of fallibilism is of course different from Cohen's (1988) and Lewis's (1996) fallibilism which says that there is *fallible knowledge*, knowledge despite *uneliminated* possibilities of error. In contrast to Lewis's contextualist model, the model we are going to talk about in this and the next section cannot be purely "eliminativist" in nature.
- ¹⁵ A similar definition in terms of belief contractions was first suggested by Gärdenfors (1988, p. 88).
- ¹⁶ As suggested ("unofficially") from a belief revision perspective by Segerberg (1998). Also cf. Segerberg (1999, p. 345). That knowledge requires maximal justification or certainty has of course been a central claim in much traditional epistemology.
- ¹⁷ In symbols: $B^*(\text{might-}\alpha) = B \cap B^*\alpha$. Likewise, a belief contraction with respect to $\neg\alpha$ can be defined by the equation $B - \neg\alpha = B \cap B^*\alpha$ which is known in the belief revision literature as the Harper identity (Gärdenfors, 1988, p. 70).
- ¹⁸ Fogelin (1994, chapter 2) and Williams (2001, chapter 4).
- ¹⁹ For a general survey of these developments and for a discussion of the methods of conservative and moderate belief change, see Rott (2003a).
- ²⁰ Rott (2003a, pp. 131–136). This reduction of iterated revisions of belief sets to one-shot revisions crucially depends on the critic's being consistent – which is guaranteed because she only speaks the truth. For the same reason, I think that the order-independence of the revisions in question is intuitively desirable. If we look at the level of systems of spheres rather than at the level of belief sets, then no reduction of iterated to one-shot revisions is possible; at this level, it becomes manifest that the method of moderate revision invariably gives priority to more recent over less recent information. For conservative belief change, a reduction of iterated revisions to revisions by conjunctions is impossible even at the level of belief sets.
- ²¹ See Rosenthal (2001, pp. 546–547; 2003, pp. 254–255). His discussion is inspired by Lehrer's (1990, pp. 137–140; 2000, pp. 156–160) recent discussions of examples for misleading evidence, viz., the Grabit and the newspaper examples. Lehrer takes the Grabit example to refute Klein's (1971) proposal. The criterion Lehrer takes as decisive in this context is the "dependence on a false belief", but I doubt that this notion can carry the theoretical weight necessary for the separation of knowledge from mere belief.
- ²² Most contextualists assume that a subject's *knowledge* set is closed under *known* logical implication. What I am assuming here is different: That the subject's *belief* set is closed under logical implication. Of course this assumption is not realistic for explicit beliefs. But it makes good sense for implicit beliefs (beliefs ascribed or beliefs the subject is committed to). Advocates of the assumption include Daniel Dennett, Isaac Levi and Robert Stalnaker.
- ²³ Proof: Suppose S is provided with new belief-contravening information, α , and S believes $\neg\alpha$ before the revision. Suppose further that β is some truth that is not believed after the revision has taken place. Then, by the deductive closure of S 's prior and posterior belief sets, $\alpha \supset \beta$ is a true proposition that is believed before, but not after the revision (for more details, compare Observation 6 of Rott 2000). If we wanted to show that $\alpha \supset \beta$ is a piece of *knowledge* that is being lost, we would have to flesh the story out in such a way that the original belief $\alpha \supset \beta$ was not dependent on the false belief $\neg\alpha$.
- ²⁴ Motivated by some recalcitrant examples, both Nozick and DeRose instruct us that in judging similarities we need to give a lot of weight to the subject's method

or way of coming to believe α . This is not an aspect that would normally be regarded as important for determining a possible world's overall similarity with the actual world.

- ²⁵ This assumption is made just for the sake of argument. It is dubious even if first and third persons coincide. There are at least three serious problems: (i) Plausibility for the subject is conceptually different from similarity with the actual world; (ii) a sphere system with $\{w_a\}$ as its innermost sphere, epistemically interpreted, represents a situation in which the subject is both infallible and omniscient; (iii) since belief is transparent to the agent, one would probably expect that $\text{Bel}_S O$ and $\text{Bel}_S \neg H$ hold throughout all spheres from the first person perspective.
- ²⁶ Skeptical hypothesis are far-fetched in the sense that the first worlds satisfying them can be found only in the periphery of the system of spheres. This is an important presupposition of DeRose's account. Michael Williams has made it clear (in discussion) that he strongly disagrees with this view.
- ²⁷ A recent paper on "Belief revision and epistemology" by Pollock and Gillies (2000) does not address the same problem. It rather compares the very special system of nonmonotonic reasoning invented by Pollock with standard belief revision theory, notes that the two do not fit together and puts all the blame on the latter theory. I think the divergence is easily explained using the terminology suggested in Rott (2001, chapter 3): Pollock takes the "vertical perspective" while standard belief revision theory takes the "horizontal perspective", and one cannot take two perspectives at the same time.

REFERENCES

- Cohen, S.: 1988, 'How to Be a Fallibilist', *Philosophical Perspectives* **2**, 91–123.
- DeRose, K.: 1995, 'Solving the Skeptical Problem', *Philosophical Review* **104**, 1–52. (Page references to the reprint in Sosa and Kim (eds.), (2000), 482–502.)
- Descartes, R.: 1641/1984, 'Author's Replies to the Second Set of Objections', in Cottingham J., Stoothoff R. and Murdoch D. (eds.), *The Philosophical Writings of Descartes*, Vol. 2, Cambridge University Press, Cambridge, 93–120.
- Fogelin, R.: 1994, *Pyrrhonian Reflections on Knowledge and Justification*, Oxford University Press, Oxford.
- Gärdenfors, P.: 1988, *Knowledge in Flux*, MIT Press, Cambridge, MA.
- Gärdenfors, P.: 1990, 'The Dynamics of Knowledge and Belief: Foundational vs. Coherence Theories', *Revue Internationale de Philosophie* **44**, 24–46.
- Gettier, E.: 1963, 'Is Justified True Belief Knowledge?', *Analysis* **23**, 121–123.
- Goldman, A. H.: 1987, 'Nozick on Knowledge: Finding the Right Connection', in Luper-Foy S. (ed.), *The Possibility of Knowledge: Nozick and His Critics*, Rowman and Littlefield, Totowa, NJ, 182–196.
- Grove, A.: 1988, 'Two Modellings for Theory Change', *Journal of Philosophical Logic* **17**, 157–170.
- Harman: 1986, *Change in View*, MIT Press, Cambridge, MA.
- Klein, P.: 1971, 'A Proposed Definition of Propositional Knowledge', *Journal of Philosophy* **68**, 471–482. (Page reference to the reprint in Sosa and Kim (eds.), 2000, 60–66.)
- Lehrer, K.: 1965, 'Knowledge, Truth, and Evidence', *Analysis* **25**, 168–175.
- Lehrer, K.: 1990, *Theory of Knowledge*, Routledge, London.

- Lehrer, K.: 2000, *Theory of Knowledge*, second, substantially revised edition, Westview Press, Boulder.
- Lehrer, K.: 2003, 'Coherence, Circularity and Consistency: Lehrer Replies', in Olsson, E. (ed.), *The Epistemology of Keith Lehrer*, Kluwer, Dordrecht, 309–356.
- Lehrer, K. and Paxson, T. Jr.: 1969, 'Knowledge: Undefeated Justified True Belief', *Journal of Philosophy* **66**, 225–237.
- Levi, I.: 1980, *The Enterprise of Knowledge*, MIT Press, Cambridge, MA.
- Levi, I.: 2004, *Mild Contraction: Evaluating Loss of Information Due to Loss of Belief*, Oxford University Press, Oxford.
- Lewis, D.: 1973, *Counterfactuals*, Blackwell, Oxford.
- Lewis, D.: 1996, 'Elusive Knowledge', *Australasian Journal of Philosophy* **74**, 549–567. (Page reference to the reprint in D.L., *Papers in Metaphysics and Epistemology*, Cambridge University Press, Cambridge, 1999, 418–445.)
- Lindström, S. and W. Rabinowicz: 1991, 'Epistemic Entrenchment with Incompatibilities and Relational Belief Revision', in Fuhrmann A. and Morreau M. (eds.), *The Logic of Theory Change*, LNAI 465, Springer, Berlin, 93–126.
- Luper-Foy, S. (ed.): 1987, *The Possibility of Knowledge: Nozick and His Critics*, Rowman and Littlefield, Totowa, NJ.
- Makinson, D.: 1997, 'On the Force of Some Apparent Counterexamples to Recovery', in Valdés E. G., et al. (eds.), *Normative Systems in Legal and Moral Theory: Festschrift for Carlos Alchourrón and Eugenio Bulygin*, Duncker and Humblot, Berlin, pp. 475–481.
- Nozick, R.: 1981, *Philosophical Explanations*, Belknap Press, Harvard University Press, Cambridge, MA.
- Plato: 1994, *Meno*, in *Plato's Meno in Focus*, transl. and ed. J. M. Day, Routledge, London.
- Pollock, J. L. and A. S. Gillies: 2000, 'Belief Revision and Epistemology', *Synthese* **122**, 69–92.
- Rosenthal, J.: 2001, 'Einige Bemerkungen zum Gettier-Problem', *Zeitschrift für philosophische Forschung* **55**, 540–555.
- Rosenthal, J.: 2003, 'On Lehrer's Solution to the Gettier Problem', in Olsson E. (ed.), *The Epistemology of Keith Lehrer*, Kluwer, Dordrecht, 253–259.
- Rott, H.: 2000, 'Two Dogmas of Belief Revision', *Journal of Philosophy* **97**, 503–522.
- Rott, H.: 2001, *Change, Choice and Inference*, Oxford University Press, Oxford.
- Rott, H.: 2003a, 'Coherence and Conservatism in the Dynamics of Belief. Part II: Iterated Belief Change Without Dispositional Coherence', *Journal of Logic and Computation* **13**, 111–145.
- Rott, H.: 2003b, 'Lehrer's Dynamic Theory of Knowledge', in Olsson E. (ed.), *The Epistemology of Keith Lehrer*, Kluwer, Dordrecht, 219–242.
- Seegerberg, K.: 1998, 'Irrevocable Belief Revision in Dynamic Doxastic Logic', *Notre Dame Journal of Formal Logic* **39**, 287–306.
- Seegerberg, K.: 1999, 'Default Logic as Dynamic Doxastic Logic', *Erkenntnis* **5**, 333–352.
- Shope, R. K.: 1983, *The Analysis of Knowing – A Decade of Research*, Princeton University Press, Princeton NJ.
- Sosa, E. and Kim J. (eds.): 2000, *Epistemology – An Anthology*, Blackwell, Oxford.
- Spohn, W.: 1988, 'Ordinal Conditional Functions', in Harper W. L. and Skyrms B. (eds.), *Causation in Decision, Belief Change, and Statistics*, Vol. II, Reidel, Dordrecht, 105–134.

- Spohn, W.: 2003, 'Lehrer Meets Ranking Theory', in Olsson E. (ed.), *The Epistemology of Keith Lehrer*, Kluwer, Dordrecht, 129–142.
- Williams, M.: 2001, *Problems of Knowledge*, Oxford University Press, Oxford.

Department of Philosophy
University of Regensburg
93040 Regensburg
Germany
E-mail: hans.rott@psk.uni-regensburg.de

LYDIA MECHTENBERG

THE STABILITY THEORY OF KNOWLEDGE AND BELIEF REVISION: COMMENTS ON ROTT

ABSTRACT. In this commentary on Rott's paper "Stability, Strength and Sensitivity: Converting Belief into Knowledge", I discuss two problems of the stability theory of knowledge which are pointed out by Rott. I conclude that these problems offer no reason for rejecting the stability theory, but might be grounds for deviating from the standard AGM account of belief revision which Rott presupposes.

1. INTRODUCTION

In his very inspiring paper, Rott explores the possibilities of reformulating within the framework of the AGM style belief revision model¹ theories of knowledge which are not based on the concept of justification. The criteria belief stability, strength of belief, sensitivity for truth and strength of epistemological position – all thought to capture at least some aspects of knowledge – are explicated within a model of spheres of possible worlds which is founded on AGM style theories of one-shot belief revision.

Thereby, Rott makes explicit the interconnections between belief revision theory and some of the mainstream epistemological criteria for knowledge in the post-Gettier age. Moreover, armed with the tools of a formal theory, which, however, are sufficiently kept in the background, Rott is able to clarify many of the conceptual and logical relations between the above-mentioned criteria for knowledge. Every epistemologist will be grateful to him for this.

The major part of Rott's paper is concerned with the so-called stability theory of knowledge, hereafter called ST. The ST is meant to capture the intuition that any belief which counts as knowledge



should not be given up in reaction to new, true, incoming information. Among the criteria for knowledge already mentioned, the stability criterion is the one which is most critically discussed by Rott, and which opens up a systematic philosophical discussion. There are two problems with the ST which Rott diagnoses.

The first problem occurs in the case where the epistemic subject in question only has beliefs which are true. As Rott makes clear by means of a very elegant proof, “if all of the subject’s beliefs are true, each and every belief of his constitutes a piece of knowledge”. But according to Rott, this result is counterintuitive: “Intuitively, having only true beliefs does not protect *S* [the epistemic subject; L.M.] against being dissuaded from believing a particular one of his beliefs”. (Rott, 2004, p. 477)

The second problem with the ST is well-known. In the case where the subject in question also has some false opinions, it is surprisingly easy for a truthful critic to talk him out of some of his true beliefs. As Rott points out: “One well-entrenched false belief erases as it were a lot of putative knowledge (...)”. (Rott, 2004, p. 482) This is counterintuitive as well: “The mere fact that the subject has a false belief *q* that is sufficiently well-entrenched to drive out the true belief *p* should not in itself be sufficient to discredit *p*’s claim to the status of knowledge”. (Rott, 2004, p. 483)

Rott himself does not solve these problems for the ST; but neither does he regard them as sufficient for rejecting the ST. He explicitly leaves the question open. It is a central question, central also to Rott’s approach to post-Gettier epistemology: Whether or not one could deal with the problems of the ST within the framework of AGM belief revision theories might also be a test question for the AGM belief revision theories. Dealing with this open question is therefore the aim of my commentary. I will discuss the two problems of the ST sequentially, with the main focus on the second one, hereafter called the “stability problem”. My conclusion will be that the ST should not be rejected due to the two problems mentioned by Rott. Instead, there might be reasons for deviating from the standard AGM belief revision theories.

My commentary proceeds as follows. The second section includes a short description of the stability problem and the way in which it arises within the AGM theory of belief revision. In the rest of the paper, I will propose and partly defend a contextual account of believing which could free the ST from the stability problem. The implications of my considerations for Rott’s approach to the ST are summarized in a short concluding section.

2. THE STABILITY PROBLEM

The ST says that a current belief is a piece of knowledge only if it cannot be eliminated in the course of a belief revision process by the help of true but misleading information. Suppose that there is a sequential game G played by one player and a kind of truth machine. At each stage of the game, the truth machine reveals to the player a true sentence or proposition φ which, by the rules of the game, the player is forced to integrate into his belief-set. Suppose now that at the stage t of the game, the sentence or proposition ψ has been a piece of the player's knowledge and therefore a part of his belief-set. Then the ST demands that ψ cannot be eliminated by any action of the truth machine at any stage $T \geq t$ of the game G . New information, if true, cannot deprive us of our old knowledge – indeed this is what we should wish at least. But why should the stability condition be implemented into a full-bodied theory of belief revision?

The answer is easy. According to the ST, a belief is a piece of knowledge only if it stays stable during the sequential game G , surviving each piece of new true information which is forced onto the player by the truth machine. Thus, the game G is a process which leads to a revision of the player's original belief-set. Therefore, one cannot believe the ST without committing oneself to some or other belief revision theory. The question is therefore: which belief revision theory should one accept? Rott, both in his current paper as well as in his overall work, is committed to a so-called AGM style model of belief revision. In his view, a model in the AGM style would be the right foundation of the ST.

So where does the stability problem sneak in, and what protects subjects from it? The Preservation Axiom of the AGM theory of belief revision² makes sure that the individual S will give up any of her old beliefs only in the case of a logical conflict between the new piece of information and some of those old beliefs. Thus, the stability problem can arise only if the new true piece of information forced onto the player at the stage t of the game G contradicts some of the player's old beliefs. This fact already implies that a player who has only true beliefs is perfectly safe from any stability problems which might occur in the game G . The truth machine only confronts S with true pieces of information, so that no contradiction can arise between these inputs and S 's old beliefs. Having only true beliefs does indeed shelter subjects from the stability problem, as Rott points out – but it does so if and only if the ST is reformulated in a framework in which

only logical conflicts between the new inputs and the old belief-set result in belief revision. Indeed, this approach does not seem to capture all of those intuitions which are comprehended by the notion of stable belief. Intuitively, a stable belief should be a conviction which we do not give up in reaction to a new true informational input – not if this input contradicts our belief, but also not if this input makes our belief highly improbable or implausible without contradicting it. If we take into account this last-mentioned aspect, then the ST allows for the stability problem arising even if the epistemic subject has only true beliefs to begin with. Therefore, the first problem which Rott attributes to the ST – the fact that a subject having only true beliefs is *ipso facto* safe from the stability problem – should not be attributed to the ST in itself, but to the somehow limited framework of at least some belief revision theories which cannot cope with the full intuitive meaning of our notion of belief stability.

A completely different treatment is needed with regard to the second problem of the ST that Rott mentions, namely the stability problem. The case can be illustrated by the well-known Grabit-example which Rott has, in a slightly varied version, used in his paper. Call *S*'s belief that Grabit has stolen the book '*A*' and her belief that the book has been stolen at 3 p.m. '*B*'. In reality, *A* is true and *B* is wrong. *A* is highly and rightly justified by *S*'s perception, so that according to normal epistemological criteria one would have to say that *A* is a piece of *S*'s knowledge. *B*, though wrong, is even more justified than *A*, and thus, *B* is more "entrenched" than *A*, that is to say that the individual *S* prefers *B* over *A*: $B > A$ is assumed.³

Taken together, *A* and *B* imply that Grabit has not been at the restaurant at 3 p.m., that is to say $(A \wedge B) \rightarrow \neg C$, where *C* stands for "Grabit has been at the restaurant at 3 p.m.", so that, besides $(A \wedge B) \rightarrow \neg C$, $\neg C$ and $\neg(A \wedge B \wedge C)$ are parts of *S*'s belief-set. Unfortunately, the truth machine convinces *S* that *C* is true; and *S* has to revise her belief-set. The stability problem arises if this revision takes the following form:

C becomes a potential belief and $\neg C$ a potential non-belief.

C and $\neg(A \wedge B \wedge C)$ imply $\neg(A \wedge B)$, so $\neg(A \wedge B)$ becomes a potential belief.

S asks herself whether she should give up *A* or *B*.

A becomes a potential non-belief, because *B* is better entrenched than *A*.

All potential beliefs become real beliefs; and all potential non-beliefs become real non-beliefs. The revision process is completed.

Obviously, the third step is the crucial one. There, entrenchment relations come in, and the belief A – which we thought to be a piece of S 's knowledge – is lost.

It should be questioned whether the belief revision behaviour of S as described above is really rational. On the one hand, the choice between A and B seems to be unavoidable. If S wants to have a consistent belief-set – which she should want to have, if she is rational – she cannot believe A , B , and $\neg(A \wedge B)$ at once. On the other hand, S knows that she runs the risk of losing true and justified beliefs. How can this conflict between knowledge preservation and consistency be solved?

Intuitively, consistency means that S is not allowed to believe the conflicting sentences A and B at once. Whenever she believes A , she must not believe B ; and whenever she believes B , she must not believe A . But this rule does not entail that S must not believe A and believe B . S is allowed to believe A and to believe B as long as any context in which she believes A is different from any context in which she believes B . Thus, if and only if beliefs can be rationally relativized to different contexts, then the following strategy of belief revision would perhaps be “more rational” than the one described above:

C becomes a potential belief and $\neg C$ a potential non-belief.

C and $\neg(A \wedge B \wedge C)$ imply $\neg(A \wedge B)$, so $\neg(A \wedge B)$ becomes a potential belief.

S separates A and B . She adopts the rule that A must be believed in different contexts than B .

S adopts a searching or at least an open attitude towards each piece of information which could help to decide between A and B .

Suppose that this contextual approach to believing – which, of course, needs to be defended and fleshed out – is correct. Indeed, a formal belief revision theory does already exist among the alternatives to the AGM theories which takes into account this kind of contextualism, namely the so-called *Controlled Revision* theory.⁴ It might therefore be legitimate for the moment to presuppose that the contextual approach to believing can be justified.

Under these conditions, true but misleading information may eliminate some true beliefs like A within certain contexts. But there may be other contexts in which these beliefs are not eliminated but are still believed, whereas other beliefs – like B – are given up. Thus, it

seems as if a contextual account of believing, if it is defensible and if it will be integrated into the ST, will at least weaken the stability problem.

Consequently, the crucial question which we have to deal with is whether and how the contextual account of belief can be defended. In the following part of my commentary, I shall present two independent reasons for accepting contextualism with regard to beliefs.

3. THE CONTEXTUAL ACCOUNT OF BELIEVING AND THE ORDER-INDEPENDENCE OF BELIEF REVISION

As Rott admits⁵, our intuitions want to have it that the revision of our belief-sets in reaction to new, incoming, consistent information should, in the end, be independent of the order in which the new pieces of information are processed by our minds. Rott points out that the corresponding revision technique is moderate belief change (see Rott, 2004, p. 481). But intuitions are not enough. Moderate belief change is in need of a good epistemological foundation. A contextualist account of believing could fill in this gap.

In order to see this point, suppose that *S* plays the belief revision game *G*, and that *S* is the one who has observed Grabit stealing the book. *A*, *B*, and *C* are defined as above. The truth machine convinces *S* that *C* is true. Until now, *S* has believed $(A \wedge B)$, $\neg(A \wedge B \wedge C)$ and, consequently, $\neg C$. Now she has to conclude that $\neg(A \wedge B)$. Is it, under this condition, rational for *S* to give up *A* in order to integrate *C*?

Suppose that *S*, in the game *G*, gets first to know that *C* and then is told that $\neg B$ and adopts the following belief revision strategy:

C is a new piece of information which *S* wants to integrate into her belief-set.

C becomes a belief and $\neg C$ a non-belief.

C and $\neg(A \wedge B \wedge C)$ imply $\neg(A \wedge B)$, so $\neg(A \wedge B)$ becomes a belief.

S asks herself whether she should give up *A* or *B*.

A becomes a non-belief, because *B* is entrenched better than *A*.

Now *S* gets to know that $\neg B$ and wants to integrate this. Thus, *B* becomes a non-belief and $\neg B$ a belief.

The belief revision process is completed.

Obviously, *S* has given up her former belief *A* although *A* is true. According to the ST, one would have to conclude that *A* has not been a piece of knowledge – a conclusion which we would like to avoid. *S*

would not have given A up if she had known earlier that her belief B is wrong. Then, her belief revision strategy would have looked like this:

S gets to know that $\neg B$ and wants to integrate this new piece of information.

Therefore, $\neg B$ becomes a belief and B a non-belief.

Now the truth machine tells S that C is true.

Therefore, C becomes a belief and $\neg C$ a non-belief.

The belief revision process is completed.

Clearly, the second process of belief revision does not give rise to the stability problem because in this particular example, it does not rely on any entrenchment relations at all. Because luckily, the new pieces of information get available in the “right order”, and because there is enough informational input, pure logic suffices to establish the posterior belief-set. What can be generalized from these particular examples is the following result: As long as entrenchment relations are allowed to fully determine the decision between conflicting beliefs, and as long as we do not have an additional, *ad hoc* axiom of moderate belief change, the outcome of the process of integrating a given chain of new pieces of information depends on the order of integration. If and only if the new pieces of information are revealed in the “right order”, then no belief will be given up which would have been kept if the order of integrating the new information chain would have been different.

This result mirrors the observation made above that postponing the decision between conflicting beliefs may pay off: If the relevant piece of information will come in later, then one can make a better decision in future than one could make now.⁶ Therefore, adopting a contextualist account of believing which allows us to postpone decisions between conflicting beliefs by assigning them to different contexts would provide the epistemological justification of moderate belief change and should therefore be taken into consideration by Rott.

4. THE CONTEXTUAL ACCOUNT OF BELIEVING AND A RISK THEORETICAL DEFINITION OF “BELIEVING”

Existing belief revision theories share a common trait: They might be used as part of an action theory, although none of them in fact *includes* any kind of action theory. Our beliefs can explain our

actions, but nobody within the belief revision literature seems to think that our actions could explain our beliefs as well. I think that this is an important oversight. In order to show what I mean, it is helpful to ask what “believing” means. One answer given by epistemology is this: *S* believes that *p* if and only if in case she were asked whether she believes *p* and in case she would be disposed to answer sincerely, she would say “Yes”. Now what does it mean to be sincere?

Suppose that *S* does not intend to lie when she replies that she believes Grabit to be innocent. Nevertheless, she does not leave him alone with her books. Nor does she defend him against accusations. Nor does she take the smallest risk associated with any action, the motivating reason of which would be the belief that Grabit is innocent. Would we believe her, then, when she says that she believes Grabit to be innocent? I suppose not. I suppose we would say that even if she wants to be sincere with us, she is less than truthful with herself in claiming that she believes Grabit to be innocent. She is not sincere with us, because she is not sincere with herself. She does not really believe Grabit to be innocent.

The case could be made more difficult. For example, it could be that *S* defends Grabit against accusations and leaves him alone with some of her cheaper books, but that she would never leave him alone with some of her more valuable books after the day on which she has observed him in the library. What should we say now?

Consider another example cited from Keith DeRose (1992), p. 913, who used this example for other purposes:

Bank Case A. My wife and I are driving home on Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, “Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays”. I reply, “No, I know it’ll be open. I was just there 2 weeks ago on Saturday. It’s open until noon”.

Bank Case B. My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and very important check. If our paychecks are not

deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a very bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, "Banks do change their hours. Do you know the bank will be open tomorrow?" Remaining as confident as I was before that the bank will be open then, still, I reply, "Well, no. I would better go in and make sure".

Both bank cases have a common structure. In both cases, there is a positive probability that the bank will be open on Saturday and a positive probability that it will not. In fact, the probabilities are the same in Cases A and B. Also, there is a gain – some positive utility – if the checks can be deposited on Saturday, given that they would not be deposited on Friday. This gain seems to be the same in both cases. What differs is the loss – the negative utility – associated with the situation where the checks cannot be deposited on Saturday, given that they are not deposited on Friday. In Bank Case B, believing that the bank will be open on Saturday is much more risky than in Bank Case A. The loss which would be realized if the belief is false is larger in Bank Case B.

Keith DeRose suggests that in both bank cases he believes the same – only he does not know the same. In Bank Case B, he does not know that the bank will be open on Saturday because the context of justification is much more demanding than in Bank Case A. I would like to make a slightly different point: In the sense of "believing something in a certain context", in Bank Case B, Keith DeRose does not even believe that the bank will be open on Saturday. He does not believe it because he is not ready to incur the risk which is actually associated with acting on this belief in the situation of Bank Case B. Of course, he does not believe the contrary either, because he does not incur any risk associated with an action motivated by the belief that the bank will be closed. Firstly, he does not act on the belief that the bank will be closed. He does not deposit his checks immediately, but decides to collect more information first. Secondly, there would be no risk associated with the possible falsity of the belief that the bank will be closed on Saturday, if he acted on this belief and deposited his checks immediately. Thus, he does not believe *anything* with regard to the question whether or not the bank will be open on Saturday.

Bank Case A, however, is different. The risk associated with the belief that the bank will be open on Saturday is less than in Bank Case B, but still there is risk. If Keith DeRose drives to the bank on Saturday morning and, contrary to his expectations, the bank is

closed, then he will regret that he has gone to the bank now instead of waiting until Monday, for example. This regret constitutes a loss associated with the falsity of the belief that the bank will be open on Saturday. If, however, the bank is open when Keith DeRose comes on Saturday morning, he can deposit his paychecks as planned, and this is the gain associated with the truth of the belief that the bank will be open on Saturday. In Bank Case A, Keith DeRose is ready to incur the risk of the belief that the bank will be open on Saturday. Therefore, this belief can be ascribed to him in Bank Case A.

The context relativity of beliefs which I have just introduced consists in a relativity of beliefs with regard to situations in which the subject can choose between different possible actions and in which the beliefs in question are risky to a certain degree. The different contexts with which the ascription of a given belief to the subject might vary are characterized by the degree of risk associated with the belief in question and by the relevant possible actions of the subject. The context dependent belief in question can be ascribed to the subject *in a given context* if and only if, in this context, the belief in question is risky to a certain degree, and the subject is ready to incur the risk which is associated with the belief in the given context.⁷

What I have said just now may not accord with our intuitions. Intuitively, one would probably say that in both bank cases, Keith DeRose believes the bank to be open on Saturday, but that the degree to which he believes this is less in Bank Case B than in Bank Case A. But this intuitive interpretation of the bank cases has its problems. Firstly, it is in conflict with another intuition, namely that the readiness to incur a risk on account of one's beliefs is necessary for believing something. Secondly, it is not clear what we should understand under a belief's degree in the present context. If, for example, it is the probability of the belief's being true, then the degree is the same in Bank Case A and Bank Case B. At least with regard to those beliefs that in some empirically possible situation might become the motivating reason for some action, I propose the following definition of "believing something in a certain context", that is to say of believing*:

Let there be a context or situation Ω which is characterized by a proposition p the acceptance of which by the subject S is risky for S to a certain degree, given S 's relevant possibilities of acting and the subjective probabilities of p 's truth and falsity. That the acceptance of p is risky means that making p the motivating cause of action in Ω would, for S , result in some loss in case of p 's being false and in some gain in case of p 's being true. Then, in the context or situation Ω , S

believes* that p if and only if S is ready to make p the motivating cause of her action.

Obviously, the possible gains and losses associated with p 's respective truth or falsity vary over the range of different contexts or situations. Therefore, what one believes* also varies with the context, that is to say with the risk associated with p 's truth value. What, then, does it mean to believe something generally? What is the meaning of "believing"? I suggest the following definition:

S believes that p , if and only if in at least one empirically possible context, S believes* that p .

If this twofold definition of "believing" is acceptable, then the contextualist account of believing would turn out to be correct. Of course, the risks associated with actions motivated by a certain belief vary from context to context. Thus, what one believes* – and therefore the entrenchment structure of one's beliefs* – also varies with the context. One of the consequences is that, in all probability, the stability problem will only arise in certain contexts, but not in all. Therefore, it will very probably not arise as a general problem. The contextualist account of believing is able to shelter the ST from the stability problem.

5. CONCLUSION

In my commentary on Rott's paper, I have argued that the two problems of the ST which Rott has outlined are no reasons for rejecting the ST, but might be reasons for deviating from standard AGM theories of belief revision. The first problem consists in the counterintuitive fact that possessing only true beliefs does already protect the epistemic subject from the stability problem. This disadvantage could be dealt with by a belief revision theory in which not only logical conflicts, but also grades of improbability or implausibility could trigger belief contraction. The second problem, the stability problem itself, would become less pressing if the ST were reformulated within a belief revision theory that incorporates a contextualist account of believing.

NOTES

¹ The AGM theory of belief revision has been developed by Alchourrón, Gärdenfors and Makinson. For a very good and detailed overview see Gärdenfors and Rott (1995). See also H. Rott (2001).

- ² The Preservation Axiom says: If $\neg\varphi \notin K$, then $K \subseteq K * \varphi$. φ stands for a new piece of information, K represents the initial belief-set, and $*$ represents the function of belief revision.
- ³ I do not claim that entrenchment always go hand in hand with justification. Intuitively, a belief which is justified better than other beliefs should also be more entrenched; but nothing hinges on this assumption.
- ⁴ See D. M. Gabbay et al., (2003). I thank an anonymous referee for pointing out to me the existence of this theory.
- ⁵ See Rott (2004, note 20). If the belief revision process were not only triggered by true input sequences, but also by false ones (by $\{a, \neg a, a, \neg a\}$ or $\{a, a, \neg a, \neg a\}$, for example), it is not so clear whether or not order independence is a good thing. I thank an anonymous referee for giving me this hint. But with regard to true input sequences, I share Rott's intuitions.
- ⁶ An anonymous referee has pointed out that this idea is also at work in the *Controlled Revision* theory.
- ⁷ This account of context dependent belief might be thought to have the consequence that a proposition like the one that a bachelor is an unmarried man cannot be ascribed to anyone in any everyday context. But as long as one is ready to accept that even a definition might be false with positive (subjective) probability, this is not true. If a subject S believes a definition which is false, then S runs the risk of misunderstanding her interlocutors and being misunderstood herself, which might well cause considerable disadvantages for S . (Would you offer someone a job who does not seem to be able to employ his own language?) If, however, definitions are thought to be true with (subjective) probability 1, then, indeed, they cannot constitute context dependent beliefs in the sense defined above, because they cannot be risky. They might be language rules, rules governing our beliefs, but not beliefs themselves. Similarly, any proposition which cannot be, in everyday contexts, risky if accepted – like the proposition that today, at least one bicycle has tipped over in China – cannot constitute an everyday-context dependent belief in the sense defined above. Nevertheless, there might be extraordinary contexts where even these propositions are risky and can therefore constitute a belief in these contexts. Such contexts could occur if, for example, one is asked whether one would bet a certain amount of money on a given proposition. I thank Elke Brendel and Hans Rott for pointing out to me the need to make these consequences of my approach clear.

REFERENCES

- DeRose, K.: 1992, 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research* **52**(4), 913–929.
- Gabbay, D. M., G. Pigozzi and J. Woods: 2003, 'Controlled Revision: An Algorithmic Approach for Belief Revision', *Journal of Logic and Computation* **13**(1), 3–22.
- Gärdenfors, P. and H. Rott: 1995, 'Belief Revision', *Handbook of Logic in Artificial Intelligence and Logic Programming 4: Epistemic and Temporal Reasoning*, in D. M. Gabbay, C. J. Hogger and J. A. Robinson (eds.), Oxford University Press, Oxford, pp. 35–132.

- Rott, H.: 2001, *Change, Choice and Inference: A Study of Belief Revision and Non-monotonic Reasoning*, Oxford University Press (Oxford Logic Guides 42), Oxford.
- Rott, H.: 2004, 'Stability, Strength and Sensitivity: Converting Belief into Knowledge', *Erkenntnis* **61**, 469–493.

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

CONTRIBUTORS

Antonia Barke, Institut für Psychologie, Johann Wolfgang Goethe-Universität, Mertonstraße 17, 60054 Frankfurt, Germany; E-mail: barke@em.uni-frankfurt.de

Peter Baumann, Department of Philosophy, University of Aberdeen, Old Brewery, High Street, Aberdeen AB24 3UB, UK; E-mail: p.baumann@abdn.ac.uk

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

Stewart Cohen, Department of Philosophy, Arizona State University, P.O. Box 874102, Tempe, AZ 85287-4102, USA; E-mail: cohen@asu.edu

Wayne A. Davis, Department of Philosophy, Georgetown University New North, Washington, DC 20057, USA; E-mail: davisw@georgetown.edu

Fred Dretske, Department of Philosophy, Duke University, 201 West Duke Building, Durham, NC 22708-0743, USA; E-mail: dretske@acpub.duke.edu

Mylan Engel, Jr., Department of Philosophy, Northern Illinois University, De Kalb, IL 60115, USA; E-mail: mylan-engel@niu.edu

Gerhard Ernst, Seminar für Philosophie, Logik und Wissenschaftstheorie, Ludwig-Maximilians-Universität München, Ludwigstraße 31, 80539 München, Germany; E-mail: Gerhard.Ernst@lrz.uni-muenchen.de

Verena Gottschling, Department of Philosophy, University of Mainz, 55099 Mainz, Germany; E-mail: gottschl@uni-mainz.de

John Greco, Department of Philosophy, Fordham University, Bronx, NY 10458, USA; E-mail: greco@fordham.edu

Thomas Grundmann, Philosophisches Seminar, Universität zu Köln, Albertus-Magnus-Platz, 50923 Köln, Germany; E-mail: thomas.grundmann@uni-koeln.de

Frank Hofman, Philosophisches Seminar, Universität Tübingen, Burggasse 1, 72070 Tübingen, Germany; E-mail: f.hofmann@uni-tuebingen.de

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

Nikola Kompa, Westfälische Wilhelms-Universität Münster, Philosophisches Seminar, Domplatz 23, D-48143 Münster, Germany; E-mail: nkompa@uni-muenster.de

Dirk Koppelberg, Freie Universität Berlin, Institut für Philosophie, Habel Schwardter Allee 30, 14169 Berlin, Germany; E-mail: dirkkoppelberg@aol.com

Mark Lance, Department of Philosophy, Georgetown University, New North, Washington, DC 2005, U.S.A; E-mail: lancem@georgetown.edu

Margaret Little, Department of Philosophy, Georgetown University, New North, Washington, DC 2005, U.S.A; E-mail: littlem@georgetown.edu

Lydia Mechtenberg, Department of Philosophy, Johannes Guetenberg University, 55099 Mainz, Germany; E-mail: mechte@mail.uni-mainz.de

Hans Rott, Department of Philosophy, University of Regensburg, 93040 Regensburg, Germany; E-mail: hans.rott@psk.uni-regensburg.de

Bruce Russell, Department of Philosophy, Wayne State University, 51 W. Warren, Detroit, MI 48201, USA; E-mail: bruce.russell@wayne.edu

Gilbert Scharifi, Philosophische Fakultät I, Philosophisches Institut, Universität des Saarlandes, Postfach 151150, D66041 Saarbrücken, Germany; E-mail: Scharifi@gmx.de

Williams Michael, Department of Philosophy, The Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218, USA; E-mail: mwilliams@jhu.edu