

Disjunctivism, contextualism and the sceptical aporia

Lars Bo Gundersen

Received: 19 February 2008 / Accepted: 19 February 2008 / Published online: 10 April 2008
© Springer Science+Business Media B.V. 2008

Abstract We know things that entail things we apparently cannot come to know. This is a problem for those of us who trust that knowledge is closed under entailment. In the paper I discuss the solutions to this problem offered by epistemic disjunctivism and contextualism. The contention is that neither of these theories has the resources to deal satisfactory with the problem.

Keywords Epistemology · Scepticism · Closure · Transmission · Disjunctivism · Contextualism · Knowledge · Evidence

1 Introduction

We know a lot. Yet there are also things we do not know and even things we apparently cannot come to know. And this is how it should be. Nonetheless, this very fact troubles philosophers—at least when the things we know seem to presuppose the things we cannot come to know. More formally, what bothers philosophers are instances of the following aporia (henceforth the Aporia):

The Aporia¹

- 1) $K(o)$
- 2) $K(A) \wedge K(A \rightarrow B) \Rightarrow K(B)$
- 3) $\neg K(\neg SH)$

¹ Other interesting aporias of this sort include the following: (i) The Lottery Aporia: I know that there are three beers left in the fridge. Yet, I do not know what follows, that I am not one of those unfortunate people whose home has recently been robbed of all valuables (in my case: three beers) by a burglar. (ii) Kripke's Aporia: we know some empirical proposition p —for example, the proposition that the oven back home is switched off. Yet, we do not know its consequence q , that all evidence against p is misleading.

where ‘K(*p*)’ is short for “it is known that *p*”,² ‘*o*’ is short for some ordinary fact—for instance, that there are three beers left in the fridge—and ‘SH’ is short for some sceptical hypothesis which is incompatible with *o*, such as the hypothesis that we are all heroes of a matrix-like plot in a world devoid of beer. This is an aporia, and so any theory of knowledge that is committed to holding both (1), (2) and (3) is contradictory and ought to be rejected for that reason alone *salva rationalis*. A theory that is blatantly contradictory simply is not acceptable. The objective of this paper is to argue that the remedies offered to solve the Aporia by two influential schools of thought within contemporary epistemology—Disjunctivism and Contextualism—fall short of a satisfactory response to this challenge.

2 Knowledge and evidence

The Aporia is contradictory. Any two of its inherent claims imply the negation of the third. In other words, (at least) one of them has to give. And it better not be (2), Closure. At least, it seems to be the mainstream view in contemporary epistemology that the strategy of denying the universal validity of Closure is a ‘repugnant’ and ‘insane’ ‘non-starter’.³ This leaves us with two options: Either we must concede that we are far too *optimistic* about our epistemic situation regarding mundane states of affairs, such as how many beers are left in the fridge and, as a consequence, deny (1). Or, alternatively, we must concede that we are far too *pessimistic* about our epistemic situation when it comes to more sophisticated matters—for instance, the question whether a certain sceptical scenario obtains—and, accordingly, deny (3). But neither sort of revision is altogether happy. There is a very good reason for both our epistemic optimism and our epistemic pessimism as expressed in (1) and (3), respectively. The reason seems, furthermore, to be the *same* in either case, namely this: Knowledge is intimately linked with availability of good evidence. If, for instance, someone is said to know how many beers there are left in the fridge, then it is understood that she possesses good evidence in support of that claim. And this is just what seems to characterise the overwhelming majority of knowledge ascriptions regarding ordinary matters such as those featured in (1). Hence our optimism. But when it comes to more sophisticated matters—such as the question whether certain sceptical scenarios obtain—there just seems to be no evidence whatsoever to appeal to. Hence the pessimism.

The heart of the problem is thus what we could phrase the thesis of *availability of good evidence* (henceforth AGE):

AGE Knowledge of some subject matter *p* is closely tied up with the idea that good evidence⁴ is available for believing that *p*.

² For the present purposes it does not matter much by whom (we/I/you/people/experts/some particular subject) it is known that *p*. The only relevant constraint is that the subject(s) in question is held constant in (1)–(3).

³ This—once so popular—way out of the aporia developed by, inter alia, Dretske and Nozick nevertheless has the great advantage of also being capable of solving the aporias mentioned in footnote 1. Gundersen (2003) provides detailed and persuasive arguments for the claim that denying Closure really is a non-starter. Similar arguments are discussed in Hawthorne (2004).

⁴ In this context ‘evidence’ is taken to be neutral on the internalism/externalism issue—unless otherwise indicated.

I do not take AGE to provide any sort of conceptual analysis of the concept of knowledge in terms of necessary and sufficiency conditions. AGE should rather be understood as a good commonsensical explanation of our epistemic optimism and pessimism expressed in (1) and (3) in the Aporia.

3 Disjunctivism

Two recent influential schools of thought within epistemology have proposed interpretations of AGE that render it friendlier towards optimistic and/or pessimistic revisions of our epistemic situation and they have thus brought new hope that there might be a practicable route out of the Aporia. One school of thought—the so-called ‘disjunctive’ school—has suggested an interpretation of the ‘available’ bit in AGE according to which good evidence *is* available even for sophisticated matters such as those featured in (3).⁵ According to disjunctivism, our epistemic access to some worldly state *p* is not mediated by *p* symptoms—*p* symptoms that may also be present, and thus deceive us into believing that *p*, even when *p* is in fact not the case. Thus, for a disjunctivist, to have good evidence available is not to have the relation between a belief that *p* and *p* mediated in some favorable manner—by, say, a particularly conducive sort of ‘sensory data’. Rather good evidence consists in *being in* a certain mental state—a state of direct openness towards the relevant worldly state of affairs. Such evidence is not individuated in terms of the phenomenological upshot, which, after all, may be the same in sceptical scenarios where the subject is not at all in a state of direct openness towards the world. For the disjunctivist, evidence is rather individuated in terms of cognitive aetiology and, as such, it is not always transparent to an agent whether or not the relevant evidential state obtains.⁶ Nevertheless, when the agent comes to believe, for instance, that she is not a hero of a matrix plot (or in some other sceptical predicament)—and she comes to believe this in a tuned-in state of direct openness towards the world—then she *knows* that the relevant sceptical scenario does not obtain.⁷

Although interesting in its own light, this proposal does not provide a particularly promising solution to the Aporia. Granted, it explains why (3) in the Aporia may, after all, be false and it thus smooths out the threatening inconsistency issuing from (1) to (3). But there is an (often ignored) ‘revenge version’ of the Aporia that remains unsolved. Disjunctivism justifies optimistic revisionism regarding (3)—we may as a

⁵ McDowell (1982) laid out this programme—at least in outline. It has subsequently been revised and modified by, inter alia, Williamson (2000).

⁶ This is because the aetiology, causal and otherwise, of a given phenomenal content cannot itself feature as part of that content. Williamson has argued for some sort of semi-internalism regarding the question whether the relevant evidential state obtains. He argues that one may know (where ‘know’ is interpreted disjunctivistically) *that* the relevant evidential state obtains when it obtains. This may be so, but this should not be confused with knowledge (disjunctivistically interpreted) about *whether* the relevant evidential state obtains.

⁷ This idea immediately raises the following puzzle: By believing *p* while in the right frame of mind and thus coming to know that *p*, the agent has to be in two distinct states of mind simultaneously; a believing *p* state and a knowing *p* state. But which of these states should now be considered causally efficacious when trying to explain e.g. actions issuing from the added *p*-information?

matter of fact know that the sceptical scenario does not obtain in situations where we are tuned in in the right sort of way. But Disjunctivism does not justify any *claim* to that effect. For knowledge claims are conditional upon transparently available evidence in a way that knowledge *possession* is not. Merely knowing that the sceptical scenario does not obtain is a matter of being in the right sort of mental state. But in order to claim that one knows one must in addition be capable of *surveying* the relevant evidence upon which that knowledge claim is based.⁸ How else could one reasonably judge (and hence claim) that the evidence in question is good enough to generate knowledge? But the sort of evidence that the disjunctivist envisages we may have against a sceptical hypothesis—a mental state with at certain aetiology—is precisely a sort of evidence that is *not* survaivable since it is non-transparent whether or not the given mental state obtains. Thus, although we may know that the sceptical hypothesis is false, we are in no position to make a claim to that effect. And such a lack of knowledge *claims* regarding sceptical matters featured in (3) is inconsistent with our numerous more or less trivial knowledge claims regarding ordinary matters. This point can be made more formally with the following revenge version of the Aporia stated in terms of knowledge *claims*:

- (1*) $C(K(o))$
 (2*) $C(A) \wedge C(A \rightarrow B) \Rightarrow C(B)$
 (3*) $\neg C(K(\neg SH))$

where ' $C(p)$ ' stands for 'it is in good order to claim that p '. Note that unlike the K operator, the C operator is not factive. Often it is in good order to claim that p although p subsequently turns out to be false. Generally it is in good order to claim that p whenever one has good evidence transparently available in favour of p .

Tim Williamson has argued that it is knowledge rather than transparently available evidence that is the norm for assertion. But even so, Williamson also acknowledges a different norm for assertion defined in terms of appearance.⁹ This is the norm of 'reasonableness' (in Williamson's terminology) as opposed to the norm of 'permissibility' (likewise in Williamson's terminology), which Williamson defines in terms of knowledge. According to Williamson, it is thus in good order (reasonable) to claim that p whenever it *appears* to oneself as if one knows that p . Conversely, if all one's appearances are completely neutral on the question whether or not p , then it is not in good order (reasonable) to claim that p . My notion of 'claimability', which is defined in terms of availability of transparent evidence, thus comes very close to Williamson's notion of 'reasonableness', which is defined in terms of appearance.¹⁰ Hence, if one prefers, one can substitute the C operator in the revenge version of the Aporia above

⁸ Mutatis mutandis these considerations apply, of course, also more generally to externalist construals according to which evidence may be 'available' to a subject in a nontransparent manner.

⁹ See, for instance, Williamson (2000, p. 257).

¹⁰ More precisely, my notion of 'claimability' may be considered a generalisation of the Williamsonian notion of 'reasonableness'. Whereas 'reasonableness'—qua its definition in terms of phenomenal appearance—only finds application in cases involving perception, my notion of 'claimability' is applicable also to cases involving other epistemic sources such as testimony, memory and thinking.

with a Williamsonian R operator (where ‘R’ stands for ‘it is reasonable to assert that’).¹¹

Now, given Closure for knowledge, $K(\neg SH)$ follows from $K(o)$ and so, given Closure for the introduced C operator,¹² $C(K(\neg SH))$ follows from $C(K(o))$. But while it is in good order to claim that $K(o)$, it is absolutely not in good order to claim that $K(\neg SH)$. Once again, we thus reach a contradiction. When it comes to ordinary matters it can hardly be disputed that it appears to oneself that one has plenty of good evidence available to support one’s beliefs. But when it comes to sceptical matters no such transparent evidence will be available, not even if one happens to know—in the disjunctivistic sense of the word—that the relevant sceptical scenario does not obtain. And so, with Closure for the C operator, one gets, once again, entangled in a contradiction.

It may be granted that Disjunctivism finds the resources to deal with the original Aporia in its generous interpretation of ‘available’, according to which there is good evidence—and thus knowledge—available even for sophisticated sceptical matters like those featured in (3). But, even if *prima facie* successful, this strategy renders Disjunctivism so much the more vulnerable to the Revenge Aporia. What the disjunctivist gains by way of knowledge regarding sceptical matters, it loses by way of claimability. The disjunctivist is therefore still committed to (3*)—and, as such, the disjunctivist is still burdened by an inherent inconsistency in her theory.

4 Contextualism

Let me now turn to the interpretation of AGE offered by the second influential school of thought within epistemology, namely, contextualism. According to contextualism we ought to pay more careful attention to the ‘good’ bit (rather than the ‘available’ bit) in AGE and try to gain a better understanding of what it means to have *good* evidence available. The idea is that ‘good’, like adjectives such as ‘flat’, ‘tall’ and ‘sharp’, has no fixed and stable meaning across all contexts. What counts as sharp for a kitchen knife may be more adequately described as blunt when referring to a scalpel used for brain surgery. According to contextualism, ‘good’ in AGE should likewise be interpreted in this context-sensitive manner. In certain ‘demanding’ contexts the evidence must thus provide very strong support for the subject matter at hand in order to count as good. In other ‘easy’ contexts, however, even a rather weak piece of evidence may qualify

¹¹ Although the Closure for R would turn out equivalent to a ‘weakened’ version of Closure for C, namely: $C(K(A)) \wedge C(K(A \rightarrow B)) \Rightarrow C(K(B))$. This being so since, according to Williamson, it is reasonable to assert what, for all one’s appearances, one takes oneself to *know*—not what, for all one’s appearances, one merely takes to be the case. I owe this point to Crispin Wright.

¹² As stated in (2*). Closure for the C operator is as plausible and trivial—if not more so—as is Closure for knowledge. Like Closure for knowledge it must, though, be subjected to the following natural qualifications: (i) it must be modalised. *De facto* claimability of A does not imply *de facto* claimability of the consequent B. Rather, if one *can* put oneself in a position in which it is in good order to claim that A, then one *can* likewise put oneself in a position in which it is in good order to claim that B. (ii) It is the availability of *overall* transparent evidence that renders a proposition claimable. One may have transparent evidence available in favour of *p*—but also have transparent evidence available against *p*. This, of course, does not establish claimability of both *p* and $\neg p$.

as good.¹³ So evaluated relative to different contexts, the same piece of evidence may be both good and not good and thus both qualify and not qualify as knowledge. Thus, relative to a demanding context, we do not know sophisticated matters such as \neg SH and so (3) holds true. But neither do we know ordinary matters relative to these high standards. Although quite a reasonable amount of decent evidence might be available in support of ordinary matters, this evidence does not quite count as *good* relative to the tough standards operative in the demanding contexts, and so the evidence does not quite amount to knowledge. Hence, we regain consistency among (1)–(3) by denying (1). And likewise for the easy contexts. Relative to these contexts the sound and decent evidence available in support of ordinary knowledge claims *does* qualify as good and thus we do know everyday matters relative to the easy contexts. But relative to the easy context we also know about more sophisticated matters such as \neg SH, and so consistency is again restored—this time by denying (3).

The great virtue of contextualism is that it respects the strong intuitions underlying both (1) and (3) and, to a large extent, manages to incorporate both intuitions by relativising them to different contexts. Contextualism also has a great deal to say in favour of its pivotal thesis that ‘knowledge’ is an indexical notion and as such means different things in different contexts. Adjectives such as ‘flat’, ‘tall’, ‘sharp’ and ‘good’ *are* indexical and context sensitive in this sense, so to the extent that AGE is right in linking knowledge with *good* evidence, it must be expected that this context sensitivity will rub off on the concept of knowledge.¹⁴ Indeed, many everyday examples suggest that the truth-values of knowledge ascriptions may change according to context in just the manner predicted by contextualism.¹⁵ These virtues notwithstanding, there is also a major flaw in the proposed contextualist solution to the Aporia. The flaw is that the contextualist account of epistemic optimism, the prospect of denying (3) in easy contexts, is untenable.¹⁶ We may grant the idea that knowledge is not just context sensitive but hyper context sensitive—so that even extremely strong evidence does

¹³ Significant contributions to this contextualist movement include Cohen (1987); Lewis (1996) and DeRose (1995). One of the great challenges for contextualism is to come up with a plausible story about what it is that individuates, for instance, a ‘demanding’ context and, a fortiori, which factors may change a ‘demanding’ context to an ‘easy’ context and vice versa. In the literature on contextualism it has been proposed that the relevant factors may include (various combination of) the following: That there is much at stake (pragmatically) with the knowledge claim in question, that counter possibilities are mentioned explicitly, that epistemic terms such as ‘knowledge’ are mentioned explicitly.

¹⁴ This rather direct route to contextualism is what Jason Stanley (2005) phrases as ‘contextualism on the cheap’. In his *Knowledge and Practical Interest* he offers sophisticated and persuasive objections against certain versions of these ‘cheap’ arguments for contextualism (see in particular pp. 75–83). Stanley’s reflections do not, however, address the version of ‘contextualism on the cheap’ presented above.

¹⁵ But maybe not quite as radically as assumed in the proposed solution to the Aporia. Also, if standards really can be raised and lowered one would expect standards in some contexts to be just high enough to deprive us of knowledge of sceptical matters, yet low enough to grant us knowledge of ordinary matters. If so, the Aporia still poses a significant challenge within these contexts.

¹⁶ This is not to say that contextualism as such is an untenable position. The claim is merely that it needs to import the resources to deal adequately with the aporia from other theoretical areas. In fact, this is precisely the strategy favoured by one of the leading contextualists, Stewart Cohen. Cohen thus argues that contextualism must be supplemented by coherentist consideration in order to maintain epistemic optimism. As is hopefully clear from the concluding remarks, the present paper can also be seen as a negative defence of such a coherentist approach to the Aporia.

not count as good according to harsh standards. We may also grant, conversely, that even very poor evidence does count as good according to lowered standards. Despite this, it is still impossible to gather good evidence—and, a fortiori, knowledge—that the sceptical scenarios featured in (3) do not obtain. This is so no matter how low the standards are set. The difficulty with sceptical matters is not that the evidence at hand is too weak to bring about knowledge of them according to strict standards—but that it may qualify as knowledge if the standards are sufficiently lowered. The problem is that there is absolutely no evidence—be it good or bad—available. Lowering the standards will never do the trick of turning complete lack of evidence into knowledge.¹⁷ And even if lowering the standards in this way could, per impossible, turn water into wine, as it were, the standards in question should govern the *entire* context and *any* unsupported, true belief would thus have to qualify as knowledge in that context.

It might be objected here that lowering the standards to a suitable level *will* result in knowledge of sceptical matters since there will be some, if only very little, evidence available for the relevant sceptical matters. Does not the fact that we have not subsequently encountered two black cats recently suggest that we are not held hostage to a Matrix illusion?¹⁸ And does not the fact that our present experience is characterised by a high degree of consistency and structure indicate that we are not at present dreaming? Weak evidence, true enough, but evidence all the same and, as such, evidence that will qualify as knowledge if the standards are set sufficiently low. This proposal is reminiscent of the so-called ‘Russellian retreat’: admitting that most of our evidence, strictly speaking, does not qualify as knowledge but, all the same, serves us well enough for all practical purposes. However, the bad news is that when it comes to sceptical matters this strategy does not work. Again, this is not because weak as it is the evidence, strictly speaking, does not qualify as knowledge. The reason is that *no* such evidence is available. When there appears to be *some* evidence available, as in the two examples mentioned above, it is simply because the sceptic has not done her work properly. If she has, the scenarios she comes up with will be moulded on what Dretske once phrased as ‘contrast consequences’, or inferences of the following logical form:

$$p \rightarrow \neg(q \wedge \text{SH})$$

where q is some empirical claim incompatible with p and SH is some hypothesis relative to which any evidence for p transforms into equally strong evidence for q . Dretske’s classical illustration of a contrast consequence is given by:

(p) The animals in the pen in front are zebras.

¹⁷ Other more familiar objections include the following: Contextualism, at least when ventured as an anti-sceptical strategy, involves an implausible error theory (Stephen Schiffer). The indexicality in question does not behave like more familiar forms of indexicality despite the indexicality for epistemic terms being motivated by and moulded over these more familiar examples (Jason Stanley and John Hawthorne). The contextualist is committed to Moorean-type claims such as “I would know p relative to easy standards, and so p since knowledge is factive in any discourse; yet I do not know p in this demanding context. In short: p but I do not know that p ” (Tim Williamson and Crispin Wright).

¹⁸ Cf. the movie of the same name. Whenever the Matrix agents perform a certain operation it results in a black cat showing up twice in the victim’s illusionary environment.

It follows that it is not the case that:

(*q*) The animals in the pen in front are mules.

In particular, the animals in the pen are not mules that happen to resemble zebras because:

(SH) The zoo authorities have cleverly disguised the zoo mules to look exactly like zebras.

But similar contrast consequences can be construed for virtually any empirical proposition. *p* could, for instance, be the claim that there are three beers left in the fridge. The corresponding *q* could then be the incompatible claim that we are living in a world devoid of beer, and SH the hypothesis that a Matrix scenario has been set up to make it appear exactly as if there were three beers left in the fridge. The crucial bit is ‘exactly’ in SH. If the appearance of *q* is *exactly* like that of *p*, we would not have *any* transparent clues whatsoever to go on when trying to demarcate *q* and SH scenarios from *p* scenarios; and a fortiori absolutely no evidence available that the *q* and SH scenario does not obtain.¹⁹

Likewise it may be objected that our discussion so far of the contextualist’s understanding of evidence and, in particular, the contextualist’s notion of strength of evidence, has proceeded in terms of *reason* to believe. So, for instance, the point just stressed about the similarity in appearance of sceptical and non-sceptical scenarios was meant to motivate the claim that we cannot have any good *reasons* to suppose that one rather than the other scenario obtains. If one instead understands evidence and strength of evidence simply in terms of its capacity to exclude error scenarios within a specified range of possible worlds, then one may possess rather strong evidence for a belief in *p* although the evidence in question is perfectly compatible with *some* (SH-like) not-*p* scenarios (as long as these SH-like not-*p* scenarios are not included in the specified set of possible worlds). Indeed, this is how other leading contextualists such as David Lewis and Keith DeRose think about strength of evidence. The more possible worlds within a pre-defined order a given piece of evidence can rule out, the stronger is the evidence. On this understanding one can therefore have strong evidence in favour of some proposition *p*—evidence that rules out²⁰ a wide range of the not-*p* possibilities—and yet have no *reason* to believe that *p*. One may even have this evidence in favour of *p* without being *aware* that one has it. Such externalist versions of Contextualism may have the resources to solve the Aporia. However, as we saw in the discussion of Disjunctivism, there is a price to be paid for such an externalistic move. To grant that one can have evidence favouring and thus come to know that the sceptical hypothesis is false without having any *reason* to believe so is to let go of one significant norm regarding knowledge claims, namely, the norm of reasonableness. To be sure, this fact in itself should not trouble someone with externalist leanings. The problem is rather that the norm of reasonableness remains in place regarding knowledge claims about ordinary matters and so the Revenge Aporia will apply to these externalistic

¹⁹ Unless, of course, we choose not to individuate the relevant evidence in terms of appearance. This, however, would just take us right back to the revenge version of the Aporia discussed above.

²⁰ Ruling out in the sense that possession of the evidence is incompatible with those not-*p* possibilities.

versions of contextualism. The externalistic contextualist can establish that we possess knowledge about both ordinary matters and sceptical matters in ‘easy’ contexts. But she must admit that it is in good order to claim knowledge of ordinary matters in these contexts but absolutely not in good order to claim knowledge of sceptical matters. And so, since closure holds for claimability (in the sense previously defined) we are once again faced with an outright contradiction regarding knowledge *claims*. The externalistically minded contextualist is thus as vulnerable to the Revenge Aporia as is the disjunctivist.

It may still be objected that it *is* possible to gather some evidence—even internalistically conceived—against the sceptical hypothesis even if all appearances are neutral on this matter. For, the objection goes, although appearance is a very significant source of evidence, it is surely not the only one. One can, for instance, also gather evidence from reasoning. In particular, since it is agreed on all sides that there is good evidence available for the ordinary matters featured in (1) from which the sceptical matters featured in (3) are entailed, the obvious thing to do is, via reasoning, to *transfer* that evidence through the entailment and thus gain evidence for the corresponding sceptical matter. But, again, if the sceptic has done her work properly, pursuing this strategy will not lead us to a satisfactory solution. Intuitively, it is somehow bootstrapping thus to transfer the perceptually based evidence one has for believing that there are three beers left in the fridge to the belief that one is not caught in the Matrix. Likewise, as Dretske has pointed out, with the zebra case, it is not appropriate to transfer the perceptually based evidence one has for believing that the animal in the pen is a zebra to the belief that it is not a cleverly disguised mule. For all the creature’s distinctive zebra-like qualities, no part of all this evidence will count as evidence for the belief that it is not a cleverly disguised mule. It is, admittedly, hard to spell out precisely what this inappropriate bootstrapping comes down to. One of the most promising attempts to do so has been made by Crispin Wright. Wright’s suggestion is that the transmission of evidence simply *fails* in such cases because it is a precondition for the perceptually based evidence to have any justificatory force relative to the antecedent that one is *independently* justified in the consequent.²¹ However, no matter how the bootstrapping quality is diagnosed and spelled out, it remains a fact that one cannot gain any evidence for the sceptical matter in question by means of the transmission of evidence.²²

By way of conclusion, then, contextualism attempts to solve the Aporia by explaining how knowledge of ordinary matters and knowledge of sceptical matters come and go in tandem as the standards for knowledge are raised and lowered. This, the contextualist claims, is the mere consequence of a more sophisticated—that is, contextualistic—understanding of AGE, in particular a more sophisticated understanding of what it means to have *good* evidence available. But although raising the standards may *deprive* us of knowledge of ordinary matters, lowering the standards will never *grant* us knowledge of sceptical matters. And without this component of the theory,

²¹ See e.g. Wright (1986, 2000, 2003, 2004).

²² The universal validity of this claim has recently been contested by James Pryor (2004). But even if one follows Pryor and only ascribes it limited validity, the limited instances in which it *does* apply would still invite trouble from the Aporia.

contextualism merely amounts to the idea that knowledge of ordinary matters and knowledge of sceptical matters *go* in tandem, that is, contextualism disintegrates into scepticism.²³ Contextualism can thus, at best, deal with the Aporia by giving in to scepticism; that is, by denying (1) in all contexts.

5 Conclusion

It is the AGE thesis—the thesis that knowledge is closely tied up with the idea of the availability of good evidence—that gives the Aporia its sting. And we have seen that the two interpretations of AGE suggested by direct realism and contextualism do not bring us any closer to a satisfactory solution. Neither sophisticated readings of ‘available’ nor sophisticated readings of ‘good’—nor their combination, for that matter—pave the way for a plausible revision in the direction of epistemic optimism or pessimism. At this stage there seems to be but one option left to pursue. AGE states that knowledge is closely associated with the availability of good evidence. The attempt to gain a more sophisticated understanding of the ‘available’ and ‘evidence’ bits has led us nowhere. The key to a solution must therefore lie in a better understanding of the third and final component in the AGE thesis: the ‘evidence’ bit. I believe there is an interesting avenue to be investigated here. Maybe the sort of evidence relevant to knowledge should not be understood in terms of ‘clues’ about or ‘openness’ towards an objective and independent reality, but rather in terms of coherence in the sense of harmonizing well with an overall belief system. If so, anti-sceptical presuppositions may be justified by the very fact that they define the structure of the belief system and, as such, play a very central role, or fit extraordinarily well, in the overall belief system. However, I do not wish to elaborate further on this proposal in the present paper. My purpose here has merely been to argue that the reason why the Aporia is so recalcitrant is to be found in the AGE thesis and, in particular, that neither a direct realistic nor a contextualistic interpretation of AGE will bring us any closer to its solution.²⁴

Acknowledgements I am very pleased to contribute to this volume in honor of Crispin Wright with the present paper. Crispin’s extraordinary intellectual powers and commitment to academic philosophy have made him one of the most significant figures in contemporary philosophy. His inspiring work has been a great gift to the philosophical society. I, for one, would never have gone into philosophy had it not been for his stimulating influence.

²³ And, to be sure, it is to no avail to insist that sceptical matters *ought* not to be knowable since merely raising the question of the knowability of a context changes it into a demanding one. This may be the correct way to think about particular and explicit knowledge claims of that sort, but when articulating and discussing contextual theory we must insist on doing so on a theoretical meta-level, as it were. On this meta-level it is possible to discuss what a particular subject knows relative to both easy and tough standards—without thereby creating a tough context in which one’s ‘theory’ ought to be that no one ever knows anything. Likewise, on this theoretical meta-level, one must insist on discussing whether a subject knows sceptical matters relative to an easy context—without thereby turning the context in question into a demanding one.

²⁴ The paper has benefitted from comments and criticism from numerous philosophers. Special thanks are due to Birgit Brogaard, Eline Busck, Carrie Jenkins, Jesper Kallestrup, Joe Salerno, Asbjørn Steglick-Petersen and Crispin Wright.

References

- Cohen, S. (1987). Knowledge, context and social standards. *Synthese*, 73, 3–26.
- Davies, M. (2004). Epistemic entitlement, warrant, transmission and easy knowledge. *Aristotelean Society—Supplementary Volume*, 78, 213–245.
- DeRose, K. (1995). Solving the sceptical problem. *Philosophical Review*, 104, 1–52.
- Gundersen, L. (2003). *Dispositional theories of knowledge*. Burlington V.T., UK: Ashgate.
- Hawthorne, J. (2004). *Knowledge and lotteries*. OUP.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549–567.
- McDowell, J. (1982). Criteria, defeasibility and knowledge. *Proceedings of the British Academy*, 68, 455–479.
- Pryor, J. (2004). Is Moore's argument an example of transmission-failure? *Philosophical Perspectives*, 18.
- Stanley, J. (2005). *Knowledge and practical interest*. OUP.
- Williamson, T. (2000). *Knowledge and its limits*. OUP.
- Wright, C. (1986). Facts and certainty. *Proceedings of the British Academy*, 71, 429–472.
- Wright, C. (2000). Cogency and queastion-begging: Some reflections on McKinsey's paradox and Putnam's proof. *Philosophical Issues*, 10, 140–163.
- Wright, C. (2003). Some reflections on the acquisition of warrant by inference. In S. Nuccetelli (Ed.), *New essays on semantic externalism and self-knowledge*. MIT Press.
- Wright, C. (2004). Warrants for nothing—foundations for free. *Aristotelean Society—Supplementary Volume*, 78, 167–212.