

Dogmatism Repuzzled

Assaf Sharon and Levi Spectre

Harman and Lewis credit Kripke with having formulated a puzzle that seems to show that knowledge entails dogmatism. The puzzle is widely regarded as having been solved. In this paper we argue that this standard solution, in its various versions, addresses only a limited aspect of the puzzle and holds no promise of fully resolving it. Analyzing this failure and the proper rendering of the puzzle, it is suggested that it poses a significant challenge for the defense of epistemic closure.

0 Introduction

Gilbert Harman and David Lewis attribute to Saul Kripke¹ a puzzle that seems to saddle us with a difficult choice between giving up the principle of epistemic closure and adhering to epistemic dogmatism of a particularly objectionable kind.

“If I know that h is true, I know that any evidence against h is evidence against something that is true; so I know that such evidence is misleading. So once I know that h is true, I am in a position to disregard any future evidence that seems to tell against h .” This is paradoxical, because I am never in a position simply to disregard any future evidence though I do know a great many different things.

(Harman 1973, p. 148)

Among the different attempts to resolve the (seemingly) paradoxical consequences of knowledge, a developed version of Harman’s own proposal is widely considered successful. We shall argue that this popular solution fails. But, more importantly,

¹ Harman (1973, p. 148), Lewis (1996, p. 442). The puzzle first appeared in print in Kripke’s ‘On Two Paradoxes of Knowledge’, unpublished lecture delivered to the Cambridge Moral Sciences Club.

analyzing its failure will show that any solution sharing its premises is bound to miss the core issue invoked by the puzzle and thus shall fail to solve it.

In section 1, after presenting the puzzle at greater detail, we present and criticize Gilbert Harman's solution. In section 2 a sophisticated development of Harman's solution is presented. This is followed by yet a further development proposed by John Hawthorne in section 3 where we also present the major problem with the Harman-Sorensen-Hawthorne line. The solution, involving the notion of "junk knowledge" is shown to be inadequate. We show, first, that it does not capture all instances of the puzzle, and second, that it is tangled in difficulties of the sort commonly associated with Gettier. The puzzle of dogmatism, we argue, is an upshot of epistemic closure and therefore must be addressed at the level of generality at which this principle is endorsed.

1 The Threat of Dogmatism

Kripke's argument apparently shows that the idea that knowledge is closed under known entailment breeds an intolerable epistemic commitment of immunity to misleading evidence. Harman, however, does not identify epistemic closure as the culprit responsible for generating the apparent problem. To present his solution, let us first look at the argument in detail.

We start with the assumption that:

- 1) S knows that p is true.²

We can also quite safely assume that:

² We could start with 'S knows that p ' and proceed from this assumption to (1). To avoid nonessential objections we begin at the outset with (1).

- 2) S knows that (if p is true, then for all evidence e , if e counts against p , then e counts against a true proposition, in other words, e is misleading).³

We now present closure in a rather simplified version (nothing will turn on the simplification):

- 3) Necessarily, if S knows that p and forms the belief that q by properly deriving it from p , then S knows that q .

(1), (2) and (3) entail:

- 4) S knows that (for all evidence e , if e counts against p , then e is misleading).

Now, assume that S actually encounters evidence e' , which, she knows, supports $\sim p$. So:

- 5) S knows that e' supports $\sim p$.

(1)-(5) entail by yet another application of closure that:

- 6) S knows that e' is misleading evidence.

Knowing that it is misleading, S would be warranted in disregarding e' . Thus dogmatism is the claim that knowledge is in a sense invulnerable – once I know that p , my knowledge cannot be undermined no matter what evidence to the contrary I shall happen to come by.

³ There are different ways in which evidence can be misleading. For instance, there may be room for setting apart undermining from overriding misleading evidence (we thank Mikael Janvid for noting this). For this distinction in a different context see Pollock and Cruz (1999, p. 196) and Casullo (2003, pp. 44-5). This distinction sometimes collapses: there are cases where the known proposition is unlikely, and hence any evidence that calls the original evidence into question (undermining the evidence), automatically counts in favor of $\sim p$ (overriding evidence). Aside from these collapsed cases (which are covered by the argument) in what follows, we will use *misleading evidence* to mean evidence counting in favor of a false proposition. *Counter-evidence* and *undermining evidence* will mean evidence that counts in favor of the negation of what one believes. Undermining evidence is not dealt with here other than when it collapses into overriding evidence (i.e. it supports $\sim p$).

The way to avoid this paradoxical result, according to Harman, while preserving all of the premises, is by attending to a temporal element that is easily overlooked. (4) and (5), Harman claims, are often incompatible. The reason is the following:

Since I now know that $[p]$, I now know that any evidence that appears to indicate something else is misleading. That does not warrant me in disregarding any further evidence, since getting that further evidence can change what I know.

(Harman 1973, pp. 148-9)

In other words, at an initial time t_0 , S knows that p and therefore may also know that any counter-evidence is misleading. When, however, she encounters such counter-evidence at a later time t_1 , her knowledge of p can be undermined.⁴ With the aid of a temporal index, it becomes evident that it is not always possible for (4) and (5) to be true at the same time. This way the derivation of (6) is blocked and the dogmatist conclusion is avoided.

To make it yet more explicit, a distinction may be drawn between two types of cases.

- a. Evidence e' is strong enough to undermine S's knowledge that p (counting in favor of $\sim p$ to a sufficient degree).⁵ At t_1 , when S

⁴ The ways in which knowledge is undermined can vary. Contextualists and subject sensitive invariantists may say that encountering misleading counter evidence raises the standards of justification. Others might have other ideas. But the exact mechanism of knowledge defeat does not matter for present purposes. What does matter is that at t_0 a subject counts as knowing that evidence e' is misleading. If it turns out that there is no plausible mechanism of knowledge defeat, so much the better for our argument. Nevertheless, contextualists who do not accept the assumption that at t_0 one knows that e' is misleading are not susceptible to the arguments below. In this respect, if we are right, the comments made by Hawthorne in relation to contextualism are defused, and contextualism is restored as a contender for resolving the paradox. Nevertheless, contextualists will have to deal with the drawback of admitting that at t_0 a subject will not know that p is true (and there seems to be little temptation in doing so - see note 28). Moreover, there is a further problem, briefly sketched in footnote 29. We do not however claim to have shown that contextualist cannot find a way mitigate these challenges.

We take this opportunity to thank an anonymous referee of this journal for drawing our attention to this, as well as more thoroughgoing points which substantially improved our thinking about the puzzle.

⁵ Notice that evidence e' does not have to undermine but merely weaken S's justification for believing that p to a degree that would make S no longer count as knowing that p . It therefore only undermines the knowledge and not necessarily the justification or the evidence.

encounters e' (and no earlier), (5) is true and consequentially (1), (4) and (6) are then false.

- b. Evidence e' is either not of the right type or not of the required strength to undermine S's knowledge that p , and consequently, (4), (5) and (6) are true. S knows that e' is misleading (that it supports a falsehood).

What this shows is that S is only warranted in disregarding e' when it is not strong enough to undermine S's knowledge that p and is thus plausibly known to be misleading at t_1 (b-cases). When it does undermine S's knowledge that p , e' is not actually known to be misleading at t_1 (a-cases), and S is then not warranted in disregarding e' . In any case, epistemic dogmatism is avoided.⁶

While we find no difficulty with Harman's solution so far as it goes, we think it does not go far enough. Specifically, we will argue that this solution does not get at the root of the problem illustrated by the puzzle, and therefore does not adequately resolve it. Once the more fundamental issue is addressed, Harman's solution, ingenious though it is, is no longer required.⁷

As we have seen, Harman's solution concentrates on how to avoid the unacceptable conclusion (6). There is however a further problem regarding the acceptability of (4), i.e. the claim that, knowing that p , S can know that any counter-evidence is misleading. Independently of what is to happen at t_1 , it does not seem plausible to say that at t_0 , having only the knowledge that p is true, S can infer from the general truism that if something is true evidence against it is misleading, that counter-

⁶ Harman only considers (a)-cases, nevertheless his careful (if somewhat confusing) formulation, leaves room for (b)-cases as well. See the quote above.

⁷ It might still be required for solving a related problem having to do with justification, specifically, with the question of whether someone is justified in believing that any counter evidence is misleading. Since we do not question justification closure, we do not address this problem.

evidence to p is misleading. In other words, given the epistemic means at his disposal, the knowledge ascribed to S in (4) is dubious. While Harman provides a good account of what should hypothetically happen at t_1 , he does not address the difficulties pertaining to S's epistemic situation at t_0 . Harman's response to the paradox, that is, may avoid the unacceptable *hypothetical* consequence of dogmatism, but it does not settle the *epistemic* problem of providing a reasonable account of the state of the agent regarding knowledge. To put it differently, one way of stating the problematic upshot of the argument is, as we have seen, that knowledge is not vulnerable to future evidence. This is what it means to know that any counter-evidence to p is misleading and hence ignorable. Harman's solution, it seems, accepts this conclusion as long as such counter-evidence has not been presented. But this amounts to claiming that knowledge isn't vulnerable until it is. While it may not be a strict inconsistency, it nevertheless expresses the strain in solutions that center merely on the hypothetical consequence of the argument leaving its epistemic ramifications intact. Taking the dogmatist challenge seriously, we urge, is recognizing its *epistemic* upshots for the intolerable consequences that they are.⁸ But this is not Harman's only problem. In what follows we shall criticize an elaborately developed version of Harman's solution and show that it is vulnerable to further objections.

2. Junk Conditionals and Junk Knowledge

The question we were left with at the end of the first section is: Does S know, before encountering or being within epistemic reach⁹ of actual counter-evidence, that it is

⁸ Dan Halliday suggested the following analogy: Harman's allowance of dogmatism as long as one has not encountered counter-evidence is similar to accepting racism as long as one doesn't encounter people of a different race.

⁹ We borrow the term "epistemic reach" from Egan (2007, p. 8).

misleading? In other words, is (4) true at t_0 ? This is the epistemic question posed by the dogmatism puzzle. In this section and the next we examine two variations on Harman's solution which blunt the knowledge attribution in (4) by rendering it as merely "junk knowledge." We explain this notion and how it is meant to resolve the paradoxical implications we have described and show it to be unsatisfactory.

In developing his theory of epistemic "blindspots", Roy Sorensen (1988) provided a sophisticated version of Harman's response to the dogmatism puzzle. Sorensen claims that proposition (4) is one particular instance of a wider class of propositions, which he calls "junk-conditionals". This class of propositions is comprised of conditionals which are not "robust." A proposition q is robust with respect to information i , if and only if, given i the probability of q is high. Specifically, a conditional is robust with respect to its antecedent, if and only if, given the antecedent's truth, the probability of the conditional is not substantially lowered. Many conditionals are not robust in this sense. For instance, by learning that Mary will come to the party the probability of the following conditional will substantially drop: *If Mary will come to the party, I will dance in the street wearing my underwear on my head.* In most contexts, a conditional of this type is meant as an indirect way of asserting the negation of its antecedent. The hearer is invited, perhaps by background knowledge, context, intonation etc., to infer by *modus tollens*, rather than by *modus ponens*, what the speaker believes. In Mary's case, one will typically infer that the speaker thinks that Mary is not coming to the party rather than infer that Mary is coming and that the speaker is preparing a striptease act (which, admittedly, may be the better interpretation if Mary is, say, the speaker's strip partner).¹⁰

¹⁰ This account is based on the account of conditionals originally defended by Jackson (1979, p. 1987), according to which indicative conditionals have material implication truth conditions (what has become

Now consider the example Sorensen uses in developing his version of Harman's solution. Remembering where I had just parked my car,

7) I know that my car is in the school parking lot.

It seems undeniable that

8) I know that (if my car is in the parking lot, then if my reliable friend Doug reports otherwise, then Doug's report is misleading).

By closure it follows that

9) I know that (if Doug reports that my car is not in the parking lot, then his report is misleading).

Sorensen claims that the conditional embedded in (9) is not robust with respect to its antecedent and so (9) is what he calls "junk knowledge". Junk knowledge does not expand through *modus ponens*. That is to say that when one comes to believe that the antecedent is true (whether truly or falsely), one will give up the conditional rather than infer the truth of the consequent, or at least so one should.¹¹

Sorensen's account of this example is very convincing. It is easy to see that when I come to know that the antecedent of the conditional in (9) is true, i.e. Doug actually reports that my car is not in the parking lot, I typically give up the conditional (or at least so I should; see Hawthorne's definition below). And thus, I cannot use my knowledge of the truth of the antecedent to deduce that:

known as the *horseshoe analysis*). Indicative conditionals are true if either their antecedent is false or their consequent is true. Perhaps this is the place to mention that if in accordance with some theorist's view, indicative conditionals have no truth conditions, the problem we are considering takes on a very different form. Note, however, that there is a way to formulate the same problem using disjunction (see the discussion of Hawthorne's version of the problem below). Whether the problem in its conditional form is really different from the one in the disjunctive form is not an issue that we will take up here.

¹¹ Although (9) is true, i.e. the conditional is known, in asserting it a subject represents herself as having different evidence than she actually has, and is therefore misleading her audience. Hers is unassertible knowledge (in Jackson's technical sense).

10) Doug's report is misleading.

Sorensen's account goes beyond Harman's and relates to the epistemic question we characterized earlier. The status of (4) at t_0 (before obtaining the counter-evidence that falsifies (1)) is that of junk knowledge. The conditional *is* known, but it cannot be used to expand knowledge through *modus ponens*.

Indeed this explanation is very plausible for the case of Doug. But can it be generalized to cover all relevant cases? Our answer is "no". Sorensen's solution to the puzzle relies on the fact that if the antecedent of the conditional turns out true, I will know this. Given our normal understanding of what reporting entails, *Doug cannot report to me that my car is not in the parking lot, without me coming to know that he has so reported*. If I am not aware of his report, it seems, he has not reported it to me, but merely said it to me or *tried* (unsuccessfully) to report it to me. That is to say, in Sorensen's account one is presented with the evidence (by way of having it reported to him) and is safeguarded – as we elaborate below – from falling into error. In a sense, Sorensen's case is special; even though the conditional is not robust in respect to its antecedent, still it is guaranteed that if its antecedent is true the subject will know this. But this need not obtain for the puzzle to emerge. We shall soon consider instances of the puzzle in which counter-evidence exists while the agent isn't aware of it. Although no explicit appeal is made to this feature of Sorensen's example, it plays a role in his account, which we now turn to demonstrate.

3. Junk or Knowledge? Dogmatism Repuzzled

The following is a version of Kripke's puzzle recently presented by John Hawthorne.

Suppose there are two newspapers, *The Times* and *The Guardian*, which I trust equally well for the purposes of obtaining soccer information. With good reason: both are extremely reliable in their reporting of soccer results. I look in *The Times* and find a Manchester United victory reported. I trust the report. The report is in fact correct. Under such circumstances people are inclined to say I know both that *The Times* said that Manchester United won and also that Manchester United won. Let us suppose I also know that *The Guardian* will have reported a result for the Manchester United game. I deduce that either *The Times* and *The Guardian* correctly reported a Manchester United victory or else *The Guardian* made a mistake about the Manchester United result. Suppose, in fact that, unbeknownst to me, *The Guardian* did make such a mistake. People are not inclined to say I know the above disjunction.

(Hawthorne 2004, p. 71)

The inclination Hawthorne alludes to seems to clash with the widely acknowledged epistemic closure principle. Hawthorne (who is one of the prominent defenders of this principle) claims that the most plausible and defensible form of this principle is:

Single-Premise Closure (SPC): Necessarily, if S knows p , competently deduces q [from p], and thereby comes to believe q , while retaining knowledge of p throughout, then S knows q .

(Hawthorne 2004, p. 34)

Obviously if S knows p , S may competently deduce from p the proposition ' p or q ' by mere familiarity with the truth conditions of sentential connectives (and a bit more maneuvering will get one to the target proposition). The problem is then that the resulting knowledge generated by (SPC) from knowledge that p , conflicts with what people are "inclined to say" in such cases. They are unwilling to attribute knowledge that *either The*

Guardian reports that Manchester United won the game, or it made a mistake, to *The Times* reading soccer-fan Jim. How are we to account for the discrepancy?

Hawthorne's answer is based on Harman and Sorensen's treatment of the dogmatism puzzle. According to his account, people tend to confuse what, following Sorensen, he labels "junk disjunctive knowledge" with *useful* disjunctive knowledge. Junk disjunctive knowledge is knowledge of a proposition of which "one of [the] disjuncts is such that if one acquires a belief in its negation, one will (or at least ought to) simply throw out the disjunction" (Hawthorne p. 72).¹²

One brand of junk disjunctive knowledge (but not the only one, as Hawthorne's example illustrates) is disjunctive knowledge that is based on knowledge of one of the disjuncts. This type of junk knowledge poses no problem for (SPC), since if one comes to believe that the known disjunct is false, one loses knowledge of this disjunct and hence loses the knowledge of the disjunction. Thus, for example, if Jim learns that *The Guardian* reports something other than a Manchester United victory, he will give up the belief in the disjunction, not deduce that *The Guardian* is mistaken (or at least he ought to¹³). (SPC), therefore, does not lead to unwelcome conclusions of the type illustrated in Hawthorne's example. It is not possible to competently deduce such knowledge from *known* premises since the knowledge of the truth of one disjunct destroys the knowledge of the first premise.

It is now clear where people's inclinations go wrong, according to Hawthorne. When determining whether someone knows such disjunctions, people often confuse

¹² Hawthorne substantially improves on Sorensen's and Harman's account by providing a formulation of this principle. He candidly admits, however, that a counterfactual (or subjunctive) characterization of this sort is problematic (see his 2004, footnote on p. 72).

¹³ Assuming that one cannot know a proposition irrationally, one loses knowledge even if one irrationally holds on to the disjunction.

knowledge and useful knowledge – since knowledge of such disjunctions is not useful, people tend to think that it is not an instance of knowledge at all. In fact, Hawthorne maintains, a person can have knowledge of these disjunctions, but this knowledge is “junk”. As we can see, Hawthorne is, then, faithful to Harman’s account and Sorensen’s development of it. But there is a crucial difference between the cases they analyze. Suppose the soccer-fan-epistemologist Jim entertains the thought expressed in the following conditional:

11) If Manchester United won yesterday’s game, then if *The Guardian* reports otherwise, *The Guardian* has made a mistake, not *The Times*.

Having read *The Times*, he knows that Manchester United won yesterday’s game, and like before, we, but not Jim, know that *The Guardian* made this mistake. So with the aid of (11) and (SPC) he should be able to competently deduce (and hence know that):

12) If *The Guardian* reports that Manchester United did not win yesterday’s game, *The Guardian*, but not *The Times*, is mistaken.

As before, if we accept (SPC) we must grant that Jim knows (12). It seems to us, however, that this account overlooks the role of a false belief of the subject in forming and maintaining his belief in the truth of (12). Once this belief and its role are recognized, attributing knowledge of (12) to our hypothetical soccer fan no longer seems plausible.

The natural and straightforward rendering of the belief content of (12) is that Jim believes that if the antecedent is true, so is the consequent.¹⁴ According to the junk knowledge account, however, the belief is based on Jim’s belief that the antecedent is false. One way to see this (and we present another below) is to consider what happens

¹⁴ In fact, since (12) and Hawthorne’s disjunction are equivalent, it’s hard to see how one might resist this reading.

when the evidence, evidence on the basis of which (12) is believed, is altered. The situation is significantly different if, for instance, Jim had gone to the stadium and watched the game. Had this been the case, although he might still have the belief that *The Guardian* had not made a mistake, his belief in the truth of the conditional would not depend on it since he would believe that *The Guardian* is mistaken if he learns that it has reported a Manchester United loss. As Hawthorne's own characterization of junk disjunctive knowledge entails, had he believed the truth of the antecedent he would throw out (or ought to throw out) the entire conditional. More than just lack of belief in the truth of the antecedent, in the case we are considering he must believe in the falsity of the antecedent if he is to maintain the belief in the truth of the conditional. That is what we now turn to illustrate.

Let us return to an important difference between Hawthorne's example and Sorensen's. In the latter, we are less reluctant to attribute knowledge of the conditional since it seems that the antecedent cannot be true unless the subject is aware of it. Doug cannot report (to me) that my car is not in the parking lot without me knowing that Doug has so reported. In Hawthorne's soccer game example, however, this is not the case. *The Guardian* can (and in the example actually does) report that Manchester lost the match without Jim's knowing it has so reported. Moreover, supposing (as Hawthorne does) that *The Guardian* has made the erroneous report, the antecedent of the conditional belief is true at the time this belief is formed. Unlike the case of Doug, in which the antecedent is false at t_0 (Doug has not yet reported anything to me),¹⁵ in the case of the soccer match

¹⁵ Admittedly there is a complication here if we assume a deterministic world in which it is already determinately settled at t_0 that Doug will report to me at t_1 that the car is not in the parking lot. If the world is in fact deterministic, the antecedent of (9) is true at the outset, that is, at t_0 . Hence, I believe falsely that he will not so report. And if we agree that one cannot have knowledge that depends on a false belief, I do

the antecedent of (12) is already true at the initial time t_0 . In other words, in contrast with Sorensen's example, in the soccer match example Jim's belief in the conditional (12) depends on a false belief, namely that *The Guardian has not reported that Manchester did not win*. Hawthorne's example illustrates that the junk knowledge account needs not only to concede dependence of knowledge on a belief that *would* be false if the counter evidence materializes (which is not much of a concession because the knowledge would then be lost), but also knowledge that depends on an *actually* false belief. According to the junk knowledge account Jim's false belief is indispensable in the purported knowledge of (12): Jim is equally disposed to take the reports of *The Guardian* and of *The Times* as true. So unless he believes *The Guardian* did not report that Manchester did not win, he will not believe (and hence will not know) (12). His belief that (12) is true, although justified, depends¹⁶ on a false belief in the falsity of the antecedent of the very conditional that the junk knowledge account maintains that Jim knows.

The evidence Jim has indicates that the consequent of (12) is false (*The Guardian* is reliable and thus not likely to have made a mistake). Thus, his actual belief in the falsity of its antecedent, the very belief that allows him to justifiably believe in this conditional, is false.¹⁷ Moreover, the evidence also supports the falsity of the antecedent – if *The Guardian* and *The Times* are both known by Jim to be reliable, Jim knows there

not know (9) at t_0 . This added complication would, then, cause trouble for Sorensen's account and would further strengthen our case since there is no temptation whatsoever to say that in such cases (7) is not known at the outset.

¹⁶ The nature of this dependence is notoriously difficult to pin down, as difficult as it is to determine the proper relations between evidence and that which it supports. In light of this our strategy here is to show that, whatever the precise nature of the dependence in Gettier cases, the beliefs in the case under consideration bear it to each other.

¹⁷ False beliefs regarding the non-existence of undermining counter-evidence do not pertain to the evidence one does have for one's belief. Thus Jim can know that Manchester United won despite his false belief that evidence to the contrary does not exist. But his justification for believing the conditional relies on the false belief that the *Guardian* did not report what it actually did. So in this case, as in the Gettier cases, once the falsity of his belief is revealed one would no longer be taken to have any justification for one's belief. The false belief directly undermines Jim's knowledge, as is made evident by the exposure test below.

has been a report of the game in *Guardian* and *The Times* reported a Manchester United victory, then it is highly probable that *The Guardian* did not report otherwise. So Jim is only justified in believing that the antecedent and the consequent are false (anything else would be counter to his evidence). Hence he can justifiably believe the conditional only if he believes in the falsity of the antecedent (as well as the falsity of the consequent). On both counts Jim's beliefs are false. Thus false belief is part of Jim's actual justification for his belief in (12).

Here is a way of making the point which brings out the generality of the problem. Jim's evidence for the belief that Manchester United had won yesterday's match, is solely the report in *The Times*.¹⁸ Since he regards *The Guardian* reliable, Jim believes that

- i) If the *Guardian* reports that *r*, then it is highly probable that *r*.

Therefore, Jim believes that

- ii) If *The Guardian* reports that Manchester United did not win yesterday's game, then it is highly probable that Manchester United did not win yesterday's game.

But since knowledge entails belief, it follows from (12) that Jim believes that:

- iii) If *The Guardian* reports that Manchester United did not win yesterday's game, then *The Guardian* is mistaken (i.e. Manchester United did win the game).

¹⁸ There is a way to further simplify the case without changing any of its essential features. Say Jim read of the victory in the *Guardian* two days ago not hearing anything more of the game since then. He knows that they have reported on the game yesterday, and deduces from the knowledge of Manchester United's victory that: *If The Guardian reported yesterday that Manchester United had not won the game, then The Guardian's report from yesterday and not of two days ago is mistaken*. It is then more apparent that Jim's knowledge of this conditional depends on (i).

Assuming he is rational and holds consistent beliefs, Jim's belief in (iii) (or (12)) entails that he believes that:¹⁹

iv) *The Guardian* did not report that Manchester United did not win yesterday's game.²⁰

His belief in (iv) is part of Jim's actual justification for (iii) (and, in turn, for (12)). This is in fact precisely what the junk knowledge account asserts – a conditional belief is “junk” if it relies on the belief that the antecedent is false. But in this case that belief – namely, (iv) – is false; hence Jim relies on a false belief (at least tacitly) in order to “know” (iii) (or (12)).²¹

Although not every dependence on false belief undermines knowledge, in this case, since, as we have noted, the false belief is an indispensable part of one's actual justification, it seems that it should. This is evident if we consider the similarity with standard Gettier cases.²² Take, for instance, Russell's example regarding knowledge of the time of day based on a stopped yet luckily accurate clock.²³ In this case the subject need not explicitly entertain any false thoughts. She acquires a true belief about the time of day that relies on an implicit belief that the clock is accurate and running properly. This background belief is an indispensable part of her justification – if it is false, she has insufficient evidence for knowledge of the time of day. In the same sense Jim's belief in

¹⁹ If one believes that if p then it is highly probable that q and that if p then $not-q$, one ought in order to maximize consistency, believe that $not-p$.

²⁰ A more direct way to reach the same conclusion runs as follows: Jim justifiably believes that *The Times* and *The Guardian* are highly reliable. Jim knows that *The Times* reported a Manchester United victory, and that *The Guardian* reported the outcome of the game. Hence, Jim believes that *The Guardian* too reported a Manchester United victory. He would be irrational in believing anything else given his evidence.

²¹ Thus even if he were to infer (12) directly from the belief that Manchester won yesterday's match, Jim would still have to rely on the false belief that the antecedent is false since he believes the consequent is false. The belief in Manchester's victory itself depends on no false belief since it is based on the evidence supplied by *The Times*' report.

²² Thanks to Karl Karlander for this important point.

²³ Russell (1948, p. 154).

the conditional depends on his belief that *The Guardian* is reliable. Since his evidence supports the belief that *The Guardian* did not issue an erroneous report and since Jim in fact believes as much, his justification for believing the conditional includes the belief that *The Guardian* did not report something other than a Manchester United victory. This background assumption is likewise an indispensable part of his justification. (If anything it plays a more central role than in the Russell example since *The Guardian's* mistake is explicit in the conditional he purportedly knows.)²⁴

Some might still insist that, although Jim relies on a false belief, and one which is indeed central to his belief in (12), this reliance need not undermine Jim's knowledge of (12), as reliance on some false beliefs is hardly avoidable (e.g. that there is no misleading counter-evidence that would undermine one's knowledge) and cannot be taken to undermine knowledge.²⁵ Perhaps Jim's reliance on the false belief that *The Guardian* did not report something other than a Manchester victory is of this benign sort. But this does not seem to be the case.

We have already shown that in Jim's case the false belief plays a crucial role in his justification. To strengthen the contrast with benign false beliefs let us introduce the exposure test. We take some false beliefs to be detrimental to knowledge because when the grounds of their falsity are exposed we tend to give up the relevant knowledge ascriptions. Thus upon learning that one's belief that *either Jones owns a Ford or Smith is*

²⁴ Some might object that our argument perhaps relies on some vexed issues having to do with indicative conditionals. This, however, is not the case. Consider Hawthorne's original example. This was a case of a junk disjunctive knowledge, not junk knowledge of conditionals. In conversation John Hawthorne said that he presented this example in disjunctive form to stay clear of the well known problems having to do with indicative conditionals. We take this opportunity to thank him for his advice regarding the argument of this section.

²⁵ We may conclude that all ravens are black based on numerous observations of what appear to be black ravens, despite the fact that some of them were actually mock-ravens. Finding out that a few samples were not real ravens should not change one's belief that all ravens are black nor the justification for this belief.

in Barcelona is based on the false belief that Jones owns a Ford one gives up the belief in the disjunction (or at least one's claim to know it). Other false beliefs do not generate the same response (or expectation) upon exposure. Thus learning that a few of the samples in one's inductive evidence base were faulty does not cause one to abandon the inductive conclusion. When exposed, benign false beliefs (i.e. false beliefs that do not undermine knowledge) might require a revision of one's beliefs, but do not entail the withdrawal of knowledge ascriptions. And the reason is simple – the evidence one has is good enough to support knowledge despite the indispensability of a false belief in one's actual justification. In Jim's case, on the other hand, if he were to find out *why* his belief is false (i.e because the antecedent is in fact true) he would no longer take himself to be justified in believing the conditional and would surrender his belief in (12). Or at least so he should, since his evidence would no longer support it. His false belief, then, plays a vital role in his actual justification for (12) and is therefore not benign.²⁶ As in Gettier cases, Jim believes a truth but for all the wrong reasons.²⁷ To recap, we argued that the junk knowledge account is committed to Gettier-type beliefs counting as knowledge. To show this we made two claims. First, that beliefs the junk-knowledge account considers to be knowledge depend on a false belief. Since the agent believes that the consequent of the

²⁶ Notice that given his evidential situation Jim cannot remain neutral on these matters. Doing so would be irrational since Jim's beliefs would be grossly disproportional to his evidence. This is perhaps the place to recall Sorensen's account. The version of the problem that relates to justification and evidence is dealt with by the "junk" account for justification. Obviously when one encounters counter-evidence one can lose one's justification. Rather than challenging his account (that relies on Harman's observation) our argument relies on one having "junk" justification. Gettier's examples are, of course, meant to challenge knowledge not justification.

²⁷ We are not suggesting that we have supplied a completely general criterion distinguishing knowledge-undermining false belief and false beliefs that do not undermine knowledge. What we claim to have shown is that the present case is on a par with the central Gettier cases in which the false belief is an actual part of one's justification and not just a tacitly assumed counterfactual, and is also essential to one's justification in the sense that without it one would have none (unlike the case of a few false instances in one's inductive base, instances which would not undermine one's knowledge even if one knew the reason for which his belief is false).

conditional is false, he must rely on the antecedent's being false too. Second we claimed that this reliance on a false belief is not of the benign sort (since it is crucial to the belief in the conditional) and illustrated this by employing the exposure test.

This observation is in line with a simple, yet forceful intuition: one cannot know on the basis of one source of information that an equally reliable source would be misleading if it conveyed contrary information. The evidence one gains from the prior source may grant one knowledge of the information supplied, but discloses nothing about the veridicality of a competing source. Anyone reasoning in this way would be met with more than a hint of suspicion.²⁸

This argument, although it does not undermine Sorensen's theory for Doug-type cases, does bear on his account in two important ways. First, our analysis of the Manchester United example shows that the idea of junk knowledge cannot cover all versions of the puzzle. Sorensen can plausibly appeal to the phenomenon of junk knowledge for dealing with cases in which belief in the junk proposition does not rely on presently false belief. But, as the example of soccer-fan Jim shows, the problem can arise even when this is not the case. In cases of the latter sort it seems more plausible to claim that one does not know the conditional claim. This gives further support to the basic

²⁸ In many, perhaps all, the examples proposed as counter examples to the principle of closure, the intuition against closure seems to be based on the observation that the original item of evidence – the way in which one comes to know a proposition – does not support one of the consequences of this proposition. The same feature is in play here, yet the present case has the added advantage of being immune to the counter claim, that if one does not know the consequence on the basis of proper derivation one loses the original item of knowledge. In the present case there is no temptation to say this. We agree, then, with the standard account about the subject's epistemic state vis-a'-vis knowledge of the original item of knowledge.

Notice that contextualists need to provide some reasons for why Jim does not know the original proposition and if so how this case differs from others where knowledge is not lost. After all no false beliefs are necessarily involved in knowledge of the original proposition. In other words, even though contextualist like Lewis (1996) can explain why at t_0 Jim does not know that *if the Guardian reported a Manchester United loss, the Guardian is mistaken*, it is harder in this case to see why Jim does not know that Manchester United won the game. Perhaps the complexity of the case muddles intuitions here, but on the face of it, it is not clear that contextualists hold the key to dissolving this problem. Cf. note 29.

intuition that not only are these states not *proper* knowledge, they are not in fact items of knowledge at all. In other words, what Hawthorne takes to be a mistaken inclination on the part of people to deny such knowledge, may not be mistaken after all. But if this is true, if indeed one does not know conditional claims such as (4) and (12) (at least not on the sole basis of knowledge such as (1) and (2), or (11) – given the evidence it is based on) then Sorensen's account, insofar as it relates to the puzzle about knowledge, is superfluous. If knowledge of the junk propositions is barred, dogmatism is avoided. To put this in the terms we suggested, once the epistemic problem is solved, a solution to the hypothetical threat of epistemic dogmatism comes along for free. What we need, this is to say, is a general solution that avoids knowledge of the problematic conditionals (disjunctions, or what have you). This, however, puts considerable pressure on the principle of epistemic closure. Thus the original dogmatism puzzle is now re-puzzled as a challenge to knowledge closure.

If sound, our argument has shown significant difficulties in the central and most promising solutions to Kripke's dogmatism puzzle. The account of their failure suggests that any attempt to retract the hypothetical consequence of dogmatism while maintaining knowledge of the junk propositions from which it is derived are not only running against common intuition (as recognized even by Hawthorne), but are also entangled with Gettier-type difficulties. To fully resolve the difficulties elicited by Kripke-type arguments one must acknowledge the weight of the epistemic problem involved in attributing knowledge of the junk propositions in such cases. In essence, this line of argument basically supports the commonsensical thought that propositions of this sort cannot be known under the circumstances since the evidence available to the person does

not support them. Take Sorensen's case, for example. My memory that I just parked the car in the school lot provides proper support for my belief that the car is in the parking lot. But this evidence in no way supports the belief that Doug's report is false. Or think of Jim's case – reading in The Times that Manchester United won the match seems like proper justification for the belief that Manchester United in fact won. It is, however, utterly inappropriate evidence for knowing that if The Guardian says otherwise, it is mistaken. By insisting on the epistemic reading of the puzzle we hope to have shown that the problem of dogmatism is not one of accommodating some marginal epistemic phenomenon. Rather the puzzle is instructive regarding formal features of knowledge, in particular the closure of knowledge under entailment.²⁹ Dogmatism is thus repuzzled, this time as a direct challenge to closure.³⁰

References:

James Cargile, "Justification and Misleading Defeaters", *Analysis*, 55, 3, 216-220, 1995.

Albert Casullo, *A Priori Justification*, Oxford UP, 2003.

Stewart Cohen, "How to be a fallibilist," *Philosophical Perspectives*, 2, *Epistemology*:

²⁹ Dogmatism, we therefore believe, is on a par with the problem of "easy knowledge". The two issues involve similar problematic consequences of epistemic closure. Dogmatism in effect asks: How can one know on the basis of one's actual knowledge-supporting evidence that counter evidence will be misleading? The problem of easy knowledge can be cast as the question: How can one know on the basis of one's actual knowledge supporting evidence that this evidence is *not* misleading? Accordingly, a proper solution to any of these problems must be one that also addresses the other. We leave the full articulation of this line of argument for another occasion. Let us merely note that the close affinity between dogmatism and easy knowledge (which is a tough challenge for contextualism among other accounts of knowledge), is another reason for skepticism regarding the prospects of a distinctively contextualist solution for dogmatism. Again, although we believe this, we do not claim to have shown that contextualism is unable to resolve the puzzle. Doing so requires more extensive argument than we have space for here and would divert our present purpose. For more on "easy knowledge" see specifically: Cohen (2002, 2005), Hawthorne (2004) and Vogel (2000, 2007).

³⁰ We are indebted to Hagit benbaji, david Enoch, John Hawthorne, Dan Halliday, Mikael Janvid, Karl Karlander, Maria Lasonen-Aarnio, Krista Lawlor, Peter Pagin, and to an anonymous referee for this journal for many helpful comments. We have also benefited from responses of participants in the Philosophy of Science Seminar at Stockholm University.

- 91-123, 1988.
- Stewart Cohen, "Basic Knowledge and the Problem of Easy Knowledge," *Philosophy and Phenomenological Research* 65, 309-328, 2002.
- Stewart Cohen, "Why basic knowledge is easy knowledge," *Philosophy and Phenomenological Research*, 75(2), 417–430, 2005.
- Fred Dretske, *Perception, Knowledge and Belief*, Cambridge University Press, 2000.
- Andy Egan, "Epistemic Modals, Relativism, and Assertion," *Philosophical Studies* 133, 2007, 1–22.
- Edmund L. Gettier, "Is Justified True Belief, Knowledge?" *Analysis* 23, 1963, 121-123.
- Gilbert Harman, *Thought*, Princeton, Princeton UP, 1973.
- John Hawthorne, *Knowledge and Lotteries*, Oxford UP, 2004.
- David Lewis, "Elusive Knowledge," reprinted in his *Papers in Metaphysics and Epistemology*, Cambridge UP, 1999.
- John Pollock and Joseph Cruz, *Contemporary Theories of Knowledge*, Second Edition, Boston and Oxford: Rowman and Littlefield, 1999.
- W. V. Quine, "On a so-called Paradox," *Mind*, 62, 245, 65–7, 1953.
- Bertrand Russell, *Human Knowledge Its Scope and Limits*, Simon and Schuster, New York, 1948.
- Roy Sorensen, "Dogmatism, Junk Knowledge, and Conditionals," *Philosophical Quarterly*, 38 (October) 433– 454, 1988.
- Jonathan Vogel, "Reliabilism Leveled," *Journal of Philosophy*, 97: 602-623, 2000.
- Jonathan Vogel, "Subjunctivitis," *Philosophical Studies*, 134:73–88, 2007.