In Elia Zardini & Dylan Dodd (eds.), *Scepticism and Perceptual Justification*.
Oxford University Press (2014)

INFERENCE AND SCEPTICISM[*]

José L. Zalabardo

1. Moorean inferences

In 1939 G.E. Moore provided what he took to be a perfectly rigorous proof of the existence of the external world. It consisted in holding up his hands and saying, as he made a certain gesture with the right hand, "here is one hand" and adding, as he made a certain gesture with the left, "and here is another" (Moore 1939).

The thought that scepticism can be refuted in this way has been embraced by some contemporary epistemologists, who extract from it a strategy for dealing with challenges to our knowledge claims based on sceptical possibilities (Pryor 2004). The strategy offers an account of how we can know, contrary to what the sceptic claims, that a sceptical hypothesis doesn't obtain. The proposal is that I can know, e.g., that I am not a brain in a vat inferentially, using as (deductive) evidence any of the everyday propositions that entail this, as, e.g., the proposition that I have hands. Let me refer to an inference of this kind as a *Moorean inference*.

The claim that Moorean inferences can be genuine cases of inferential knowledge would seem to follow from a plausible sufficient condition for inferential knowledge:

> Transmission: If S knows E non-inferentially, S knows that E logically entails H and S believes H, then S knows H inferentially on the basis of evidence provided by E.[1][2][3]

[*] I have presented versions of this material at the Northern Institute of Philosophy and the University of Geneva. I am grateful to these audiences. I am also grateful for their comments to the editors of this volume and an anonymous referee.

Let's assume that I can know that envatted brains don't have hands. Then, if I can know that I have hands non-inferentially (say, by sense perception), Transmission will enable us to conclude that the Moorean inference can bestow the status of knowledge on my belief that I am not a brain in a vat.

Is this a good thing? Here we have arguments pulling in opposite directions. On the one hand, having an explanation of how we know that sceptical hypotheses don't obtain is certainly a desirable outcome. And for those who are convinced that no other explanation can work, its appeal will be almost irresistible. On the other hand, our goal in this area cannot be simply to find some hypothesis on the nature of knowledge which, if correct, would explain how we know that sceptical hypotheses don't obtain. We should seek a hypothesis that is as a matter of fact correct—one that provides an accurate representation of how we know the problematic propositions, if we do indeed know them. From this point of view, the proposal is highly counterintuitive. Our intuitions seem to be firmly against the possibility of knowing that I am not a brain in a vat by means of a Moorean inference. However, although intuitions must play an important role in the assessment of a theory of knowledge, other factors must also be taken into consideration, and we should be prepared in principle to accept a theory that yields some counterintuitive results.

The appeal of treating Moorean inferences as cases of knowledge would be drastically reduced if we could find a cogent reason for withholding this treatment from them—a feature that can be argued to pose an obstacle to the ascription of knowledge and is exhibited by

---

[1] By restricting the principle to cases in which E is non-inferentially known we rule out cases of logical circularity.

[2] It could be argued that a plausible transmission principle would need to include a clause to the effect that the subject's belief in H is based on her belief in E. I don't accept this point. See (Zalabardo 2012) for details.

[3] A subject can have both inferential and non-inferential knowledge of a given proposition. This sufficient condition for inferential knowledge might be satisfied in cases in which the subject also knows H non-inferentially.

these cases but not by unproblematic cases of inferential knowledge. If, on the contrary, we failed to identify an objectionable feature of Moorean inferences, the claim that they can produce knowledge would be harder to resist.

2. Transmission principles

A Moorean inference involves no logical circularity—its conclusion is not among its premises. Nevertheless, many have harboured the suspicion that the way in which the conclusion of a Moorean inference is related to its premise can be legitimately characterised as a form of circularity, not logical, but epistemic, and that epistemic circularity poses as serious an obstacle to inferential knowledge as the logical variety. Clearly, the challenge for this approach is to characterise epistemic circularity as a feature that is (a) present in Moorean inferences and absent from unproblematic cases of inferential knowledge and (b) intuitively incompatible with inferential knowledge.

One appealing proposal in this connection is to say that an inference is epistemically circular when the subject's knowledge of the premises requires that she has independent knowledge of the conclusion. In order to provide an adequate formulation of this proposal, it will help to have at our disposal the concept of *warrant*, understood as the property that turns a true belief into knowledge.[4] More precisely, warrant will be a relation between subjects and propositions such that S knows p just in case S believes p, S is true, and S bears to p the warrant relation.[5]

Now one constraint that is widely accepted by contemporary epistemologists is that warrant is not a primitive property—that when a proposition p has warrant for a subject S, there are some different facts in virtue of which p has warrant for S. Let me refer to the facts

---

[4] The term is used in this sense in (Plantinga 1993).

[5] Notice that the definition doesn't rule out the possibility that warrant entails belief or truth. Thus, e.g., if warrant is what Nozick calls truth tracking, warrant will entail both.

that play this role as *warrant-constituting facts*. On the assumption that warrant requires warrant-constituting facts, we can provide a formulation of the proposal under consideration. We can say that an inference is epistemically circular when the subject having warrant for the conclusion is among the warrant-constituting facts of her warrant for the premise.[6]

I think it would be hard to deny that if epistemic circularity is characterised in this way, an epistemically circular inference will be incapable of producing knowledge of its conclusion. Inferential knowledge produced by the transmission principle under consideration requires knowledge of the premises, but if an inference is epistemically circular, in this sense, the fact that the subject knows the premises will be grounded in the fact that she knows the conclusion.[7] Hence she will be incapable of using the inference to acquire knowledge of the conclusion, as the intended outcome is among the preconditions of this method of knowledge acquisition.

However, even if we accept, as I propose, that epistemically circular inferences cannot produce knowledge of their conclusions, we won't be able to use this circumstance to argue that Moorean inferences cannot produce knowledge unless we can show that Moorean inferences are epistemically circular. The issue here turns on our explanation of how we know the premise of a Moorean inference, and on many contemporary accounts of knowledge, the claim that Moorean inferences are epistemically circular has no plausibility. Suppose, for example, that I can know that I have hands by virtue of the fact that my belief to this effect tracks the truth (Nozick 1981), or that it was formed with a reliable belief-forming method (Goldman 1986). Then truth tracking, or reliable formation, will be the warrant-

---

[6] Crispin Wright (Wright 1985, 2000, 2002) has explored several versions of this diagnosis of the inadequacy of Moorean inferences. Notice that on the formulation that I am using, the presence of the warrant-constituting facts for the conclusion among the warrant-constituting facts for the premise doesn't pose a problem.

[7] There may be cases of inferential knowledge in which the premises are not known (Luzzi 2010) or not even true (Warfield 2005).

constituting fact of my belief that I have hands, and neither truth tracking nor reliable formation includes, or requires, that I have warrant for the proposition that I'm not a brain in a vat. If my belief that I have hands can obtain the status of knowledge in one of these ways, Moorean arguments are not epistemically circular.[8]

One could try to rescue this approach with a different account of why Moorean inferences are epistemically circular. The idea now would be that what's wrong with a Moorean inference is that the subject's knowledge of the premise presupposes, not that she has knowledge of the conclusion, but that the conclusion is true. We can pursue this strategy with a new account of epistemic circularity. Now an inference will be epistemically circular when the truth of the conclusion is among the warrant-constituting facts of the subject's warrant for the premise.

We could argue that on this characterisation of epistemic circularity the standard Moorean inferences are epistemically circular. The facts about me and my relationship to my environment that enable my belief that I have hands to track the truth, or to have been formed reliably, are clearly incompatible with my being a brain in a vat. If I were a brain in a vat my belief that I have hands would not track the truth, and the procedure with which I formed the belief wouldn't be reliable. Suppose it could be argued that my not being a brain in a vat is actually included among the facts that make my belief that I have hands count as tracking the truth or reliably formed. Then my inference from the premise that I have hands to the conclusion that I am not a brain in a vat would count as epistemically circular, on the construal under discussion. Then the prospects of the proposal would turn on the plausibility of claiming that inferences that exhibit epistemic circularity, on this construal, are incapable of producing knowledge of their conclusions.

[8] Versions of this point have been made in (Brown 2003) and (Pryor 2004). See also my (Zalabardo forthcoming).

It has been argued that this claim is incorrect—that there are inferences that exhibit this brand of epistemic circularity but seem perfectly capable of producing knowledge of their conclusions (Pryor 2004: 358-59; Davies 1998: 352). Here I want to argue that there is an additional source of concern for the proposal. The problem on which I want to focus is that the proposal doesn't deal with all the cases of the pathology that afflicts Moore's inference. I am going to present the argument using a process reliabilist account of non-inferential knowledge, but the argument will work in the same way for some tracking accounts.

Consider Fred Dretske's example of a child, call her Molly, who goes to the zoo and upon seeing the animals in the enclosure marked 'zebras' forms the belief that they are zebras (Dretske 1970). According to process reliabilism, Molly's belief has warrant, since it has been formed with a reliable belief-forming process. Hence, if the belief is true, if the animals are actually zebras, it will have the status of knowledge.

Suppose now that Molly knows that the proposition that the animals are zebras (call it ZEBRAS) entails the proposition that they are not mules cleverly disguised by the zoo authorities to look like zebras (call it ~MULES). It follows from our assumptions and Transmission that Molly can know ~MULES inferentially on the basis of the evidence provided by ZEBRAS. However, this inference seems to be afflicted by the same pathology as the inference from HANDS (the proposition that I have hands) to ~BIV (the proposition that I'm not a brain in a vat). It is plausible to assume that the right verdict of the difficulty that invalidates the latter inference as a case of inferential knowledge will also apply to the former.

However, on the construal of epistemic circularity currently under discussion, the inference from ZEBRAS to ~MULES would not count as epistemically circular. To see this, notice that the reliability of the process involved in the production of Molly's belief in ZEBRAS doesn't require that ~MULES is true, so long as the reliability that is required for knowledge is not perfect reliability. Molly could have formed her belief with highly reliable perceptual devices even if the animals she is looking at are cleverly disguised mules. Of course, the reliability of the relevant perceptual devices requires that this kind of deception is sufficiently rare, but not that it never occurs, or that it is not occurring on this occasion.

I want to treat the fact that the strategy doesn't deal with all the instances of the problem as an indication that its source hasn't been addressed. We haven't identified the feature of the relationship between premise and conclusion of Moorean inferences that explains why they shouldn't be treated as cases of inferential knowledge.

3. An idea from Nozick

I am going to argue that there is a more appealing explanation of why Moorean inferences shouldn't be treated as cases of knowledge. My proposal is based on the intuition that what's wrong with a Moorean inference has to do with the circumstances under which you would believe its premise. The problem is that you would still believe it if its conclusion were false. You can't have inferential knowledge of ~BIV based on the evidence provided by HANDS because envatted brains believe in HANDS, and Molly can't have inferential knowledge of ~MULES based on the evidence provided by ZEBRAS because people who are looking at cleverly disguised mules believe in ZEBRAS. Nozick formulated a principle based on this intuition, as a condition for when inferring H from E yields knowledge of H:

> I: If q were false, S wouldn't believe p (or S wouldn't infer q from p)
> (Nozick 1981: 231)

I believe this principle is along the right lines, but instead of the counterfactual formulation used by Nozick, I want to propose a formulation of the thought in terms of conditional probability. My proposal is that what's wrong with these cases is that you are not less likely to believe your evidence if the hypothesis is false than if it is true. We can reformulate the condition in terms of the notion of incremental confirmation. Evidence E confirms hypothesis

H just in case the probability of E given H is higher than the probability of E given ~H.[9] Hence we can say that the problem with a Moorean inference is that the subject's belief in its premise (that is, the fact that she believes the premise) doesn't confirm its conclusion. The proposal is, then, that we can avoid treating Moorean inferences as cases of inferential knowledge by imposing the following condition on inferential knowledge:

PI: S can have inferential knowledge of H based on the evidence provided by E only if S's belief in E confirms H.[10]

This proposal deals easily with the standard cases. I can't know ~BIV inferentially on the basis of the evidence provided by HANDS because the probability of my believing HANDS is not affected by whether I am a brain in a vat.[11] And Molly can't know ~MULES inferentially on the basis of the evidence provided by ZEBRAS because she is no less likely to believe ZEBRAS if she is looking at cleverly disguised mules than if she isn't.[12]

---

[9] This formulation of incremental confirmation is equivalent to the more standard $p(H \mid E) > p(H)$.

[10] A closer analogue of Nozick's principle, taking account of the bracketed clause, would require that H is confirmed either by S's belief in E or by her belief in the proposition that E supports H. I shall not take this line here.

[11] I think that the claim that the probability of my believing HANDS is not affected by whether I am a brain in a vat is part of what we stipulate when we describe the brain-in-a-vat scenario. If this is not, explicitly or implicitly, part of the description of the case, it might turn out that, as a matter of fact, envatted brains are less likely to believe HANDS than normal people. If this were the situation, then my inference form HANDS to ~BIV would satisfy PI. I am not sure whether, in these circumstances, it would be wrong to say that the inference produces knowledge. On this point see (Brueckner 1994: 829).

[12] A word on the notion of probability that I am assuming here. The probabilities that I have in mind are neither logical, a priori discoverable facts about events, nor subjective degrees of belief by actual or ideally rational subjects. They are instead objective, contingent facts about states of affairs, knowable only by empirical investigation. They arise from the nomological order: the probability of states of affairs is determined by the laws of nature (Lewis 1986, 1994). The notion of probability that I am assuming differs from Lewis' in two important respects. First, propositions can have non-trivial probabilities even in a deterministic world. See (Hoefer 2007; Glynn 2010) for proposals as to how to achieve this. Second, propositions about past events can have non-trivial probabilities. This is required by the thought that which evidential relations propositions bear to one another is not affected by whether or not

4. Closure and transmission

Before we proceed, we need to make a distinction between the principle that I have called Transmission and another principle that's implicated in the issues under discussion. Transmission is a principle that stipulates sufficient conditions for inferential knowledge in the specific case of deductive inferences. The other principle that I want to consider has, in the first instance, nothing to do with inferential knowledge in particular, although it also concerns cases in which E entails H. It stipulates that the following four states of affairs are incompatible: (a) S knowing E, (b) S knowing that E logically entails H, (c) S believing H and (d) S not knowing H. It is usually formulated as a conditional:

> Closure: If S knows E non-inferentially, S knows that E logically entails H and S believes H, then S knows H.[13]

In spite of the superficial similarities, there are important differences between Transmission and Closure. Notice that Transmission is stronger than Closure. On the one hand, S will know H whenever she knows H inferentially on the basis of evidence provided by E. On the other hand, S could know H in some other way. Hence Closure might be universally valid even if Transmission has counterexamples.

The most important difference between Transmission and Closure for our purposes concerns how they relate to PI. Transmission, on the one hand, is directly incompatible with PI. Whenever S knows E non-inferentially, S knows that E logically entails H and S believes H, Transmission will require that S knows H inferentially on the basis of evidence provided

---

they represent states of affairs in the past. We can achieve this by rejecting Lewis' idea that the history of the universe is taken into account in the determination of probabilities.

[13] The plausibility of the principle does not depend on the restriction to cases in which E is known non-inferentially. I introduce the restriction here to facilitate comparison with the version of Transmission under discussion.

by E. But these conditions are compatible with S's belief in E failing to confirm H, and when this happens PI will rule out inferential knowledge of H.

Hence, the adoption of PI would force us to weaken transmission as follows:

> Transmission*: If S knows E non-inferentially, S knows that E logically entails H, S believes H and S's belief in E confirms H, then S knows H inferentially on the basis of evidence provided by E.

Closure, by contrast, is not directly threatened by PI. If S's belief in E doesn't confirm H, then, according to PI, S won't know H inferentially on the basis of evidence provided by E. But this is compatible with S knowing H, by some other means, and this is all that Closure requires.

Nevertheless, the adoption of PI would also have adverse consequences for Closure. With Transmission in place, the satisfaction of Closure is guaranteed. However, as we have seen, the adoption of PI forces us to replace Transmission with Transmission*, and the latter no longer guarantees the satisfaction of Closure. If S's belief in E doesn't confirm H, Transmission* won't rule out the possibility that the antecedent of Closure is satisfied but its consequent isn't.

In sum, PI is incompatible with Transmission, but compatible both with Closure and with its negation. Hence, PI can be incorporated in a theory of knowledge for which Closure is universally valid. However, this would require ensuring that Closure is satisfied by some other means, since the account of inferential knowledge won't guarantee its satisfaction. I am not going to consider at this point whether we should take this step. Clearly PI will deprive Closure of the support that it might derive from its connection with Transmission, but the principle might be recommended by independent considerations.

5. Reflective knowledge

I want to consider next another type of inference that raises similar issues to those that we have considered in connection with Moorean inferences. Suppose that I read in a reliable newspaper that the Bulls won the game last night and form as a result the belief that the Bulls won the game.[14] On a reliabilist or truth-tracking account of knowledge, if the Bulls did indeed win, my belief to this effect will have the status of knowledge.

Consider now the proposition that the newspaper report was veridical. As we are about to see, there is an issue as to how this proposition should be analysed, but on each plausible analysis it is logically entailed by propositions that I know—the proposition that the Bulls won and the proposition that the newspaper report says that the Bulls won. There is also no reason why I shouldn't know this entailment. Assuming that I do, I will know that the proposition that the report was veridical is a logical consequence of propositions that I know. If this is sufficient for inferential knowledge, I will know that the report is veridical inferentially, on the basis of evidence provided by the proposition that the Bulls won the game.

Intuitively this is the wrong result. If the newspaper report is my only source of information for the match result, I cannot use it as evidence to obtain inferential knowledge of the veracity of the report. Notice that I am not arguing that it is impossible to know the veracity of the report or that it is possible to know that the Bulls won without knowing that the report was veridical. All I am arguing is that if knowledge that the report was veridical is to be inferential, then, in the circumstances that I have described, it cannot be based on evidence provided by the proposition that the Bulls won the match. In other words, I am arguing that the inference from the premise that the Bulls won the match to the conclusion that the newspaper report was veridical should be a counterexample to Transmission. I am

[14] Keith DeRose discusses this case in (DeRose 1995: 18).

not taking sides on the question whether we can obtain from this case a counterexample to Closure.

I want to argue that PI succeeds in ruling this out as a case of inferential knowledge. Now, whether this result holds depends on how we analyse the proposition that the newspaper report is veridical. One natural approach is to analyse it as a truth function of the proposition that the Bulls won (BULLS) and the proposition that the newspaper reported that the Bulls won (REPORT). There are at least two plausible options as to how to do this. The first is to analyse it as the proposition that the newspaper didn't falsely report a Bulls' victory, i.e. ~(REPORT & ~BULLS). The second is to treat it as the proposition that the newspaper reported a Bulls victory veridically, i.e. REPORT & BULLS.

Which of these options we take is going to make a difference to whether the inference contravenes PI. Consider first REPORT & BULLS. My evidence for this proposition will consist of the propositions REPORT and BULLS. Hence, in order to assess the inference from the point of view of PI we need to determine whether p(Bel(REPORT) & Bel(BULLS) | REPORT & BULLS) is greater than p(Bel(REPORT) & Bel(BULLS) | ~(REPORT & BULLS)). I think it's clear that the answer is yes. Notice that ~(REPORT & BULLS) is the same proposition as ~REPORT ∨ (REPORT & ~BULLS).[15] We can assume that p(Bel(REPORT) & Bel(BULLS) | REPORT & ~BULLS) is the same as p(Bel(REPORT) & Bel(BULLS)) | REPORT & BULLS), since the newspaper report is my only source of information about the game result. However, p(Bel(REPORT) & Bel(BULLS) | ~REPORT) can be expected to be much lower, since I am unlikely to believe in a nonexistent newspaper report. Hence, to show that p(Bel(REPORT) & Bel(BULLS) | REPORT & BULLS) is greater than p(Bel(REPORT) & Bel(BULLS) | ~(REPORT & BULLS)), it will suffice to show that

[15] I am assuming, for simplicity, that propositions are individuated semantically, i.e. up to logical equivalence.

p(Bel(REPORT) & Bel(BULLS) | REPORT & ~BULLS)

$>$

p(Bel(REPORT) & Bel(BULLS) | ~REPORT)

entails

p(Bel(REPORT) & Bel(BULLS) | REPORT & ~BULLS)

$>$

p(Bel(REPORT) & Bel(BULLS) | ~REPORT ∨ (REPORT & ~BULLS))

This can be easily shown (see Appendix).

I want to suggest that the way in which this result is obtained should render the analysis suspect. The reason why PI is satisfied is due entirely to the low probability of my belief in REPORT if REPORT is false. But intuitively this circumstance should not affect the adequacy of a piece of evidence as support for the veridicality claim. It seems that this should be assessed exclusively in terms of how the probability of my belief in the evidence is affected by whether or not the report is veridical. The situation in which the report doesn't exist shouldn't come into play.

Let's consider now the other proposal as to how to analyse the veridicality claim as a truth function of REPORT and BULLS, i.e. to take it as ~(REPORT & ~BULLS). Notice first of all that now REPORT is no longer needed as a premise, since ~(REPORT & ~BULLS) follows from BULLS alone.[16] Hence, in order to determine whether the inference satisfies PI on this construal, it will suffice to compare p(Bel(BULLS) | ~(REPORT & ~BULLS)) with p(Bel(BULLS) | REPORT & ~BULLS). I am going to argue that p(Bel(BULLS) |

[16] REPORT is, of course, my evidence for BULLS. The point I'm making here is that, as REPORT is ineffectual as evidence for ~(REPORT & ~BULLS).

~(REPORT & ~BULLS)) is actually smaller than p(Bel(BULLS) | REPORT & ~BULLS), contrary to what PI calls for.

Notice first that on the assumption that I have no source of information about the match result other than the newspaper report, we have that p(Bel(BULLS) | REPORT & ~BULLS) = p(Bel(BULLS) | REPORT & BULLS) = p(Bel(BULLS) | REPORT). Hence it will suffice to show that p(Bel(BULLS) | ~(REPORT & ~BULLS)) is smaller than p(Bel(BULLS) | REPORT & BULLS).

Notice that p(Bel(BULLS) | ~(REPORT & ~BULLS)) can be rewritten as p(Bel(BULLS) | ~REPORT ∨ (REPORT & BULLS)). Now, clearly, p(Bel(BULLS) | ~REPORT) is smaller than p(Bel(BULLS) | REPORT & BULLS), given my propensity to believe the report. Hence (see Appendix), we have that p(Bel(BULLS) | ~(REPORT & ~BULLS)) is smaller than p(Bel(BULLS) | REPORT & BULLS), as desired. We can conclude that Bel(BULLS) doesn't confirm ~(REPORT & ~BULLS), and hence that on this construal of the veridicality claim, the inference is ruled out by PI.

I suggested above that this is the outcome that is in line with our intuitions about this kind of case. However, this construal is open to the same objection as the previous one, since we are still taking into account the probability that I believe the evidence if the report doesn't exist. Inspection of the argument shows that p(Bel(BULLS) | ~(REPORT & ~BULLS)) is dragged down by the relatively low value of p(Bel(BULLS) | ~REPORT), i.e. by the low probability that I believe that the Bulls have won in the absence of the report. Hence the reason that I have offered for rejecting the previous proposal cannot be used as a reason for preferring this alternative. If the right analysis of the veridicality claim cannot make the admissibility of a piece of evidence depend on what I would believe if the report didn't exist, the second proposal is as inadequate as the first.[17]

I want to try a different approach to the analysis of veridicality claims. My proposal is to analyse the proposition that the newspaper report is veridical as ascribing a predicate (veridical) to an individual picked out by a definite description (the newspaper report of the Bulls' victory). Thus, if V stands for …*is veridical*, and R stands for …*is a (unique) newspaper report asserting that the Bulls won the match*, the veridicality proposition can be symbolised as V ɿx Rx. Clearly, REPORT and BULLS logically entail V ɿx Rx. Hence, according to Transmission*, I will be able to have inferential knowledge of V ɿx Rx on the basis of the evidence provided by REPORT and BULLS unless the case is ruled out by PI.

In order to determine whether this inference satisfies PI, we need to compare p(Bel(REPORT) & Bel(BULLS) | V ɿx Rx) with p(Bel(REPORT) & Bel(BULLS) | ~V ɿx Rx). PI will be satisfied just in case the former is greater than the latter. The issue turns on the familiar ambiguity of scope afflicting ~V ɿx Rx. If we take the definite description to have narrow scope, the proposition will be true if the report is not veridical, or if it doesn't exist (or if it's not unique). If we take it to have wide scope, the proposition will be true just in case the report is not veridical, i.e. just in case there exists a (unique) newspaper report asserting that the Bulls won the game and this report is not veridical.

Given the source of our dissatisfaction with previous analyses of veridicality, it should be clear that the wide-scope reading is to be preferred.[18] Adopting the narrow-scope reading

---

[17] Elia Zardini has suggested to me that these analyses of the veridicality proposition can be rejected on independent grounds. On the one hand, REPORT & BULLS obviously entails BULLS, but it could be argued that this shouldn't count as a logical consequence of the proposition that the newspaper report is veridical. On the other hand, ~(REPORT & ~BULLS) is compatible with the proposition that the newspaper falsely reported that the Lakers won, unlike the proposition that the report is veridical.

[18] Notice that, on the wide-scope reading, p(A | V ɿx Rx) > p(A | ~V ɿx Rx) is no longer equivalent to p(V ɿx Rx | A) > p(V ɿx Rx).

would assign a role in the assessment of my evidence for the veridicality proposition to how likely I am to believe the evidence if the report doesn't exist. With the wide-scope reading, however, this factor is completely excluded. Now the assessment of my evidence will depend exclusively on what I am likely to believe if the report is veridical and if it is not veridical, as intuition recommends.

Once the veridicality claim is analysed in this way, it is clear that my evidence for it doesn't satisfy PI. The probability of my believing REPORT and BULLS is unaffected by whether or not the report is veridical. This is the reason why the information that I have obtained from the report cannot be used as evidence of its veridicality.[19]

6. Not falsely believing

Similar issues are raised by beliefs concerning the truth value of my own (current) beliefs. Take, for example, my belief that I don't falsely believe A, where A is a proposition that I also believe. How could this belief acquire the status of knowledge? Here I want to discuss one possible answer to this question—the view that I can know that I don't falsely believe A inferentially on the basis of the evidence provided by A. Clearly A logically entails the proposition that I don't falsely believe A. Hence, if we assume that I know this entailment, and that I know A non-inferentially, the antecendent of Transmission will be satisfied. Therefore, whether Transmission* treats this as a case of inferential knowledge will depend on whether PI is satisfied.

I want to suggest that this inference is intuitively objectionable for exactly the same reason as the inferences that we have already considered. I argued that the inference from HANDS

---

[19] Elia Zardini has suggested to me an alternative strategy for discounting the effect of what I am likely to believe if the newspaper report doesn't exist. Zardini's proposal is to consider whether PI is satisfied on the assumption that p(REPORT) = 1. It is easy to see that, on this assumption, my inference for the conclusion that the report is veridical violates PI on both truth-functional construals of this conclusion.

to ~BIV cannot produce knowledge of its conclusion because if the conclusion were false, if I were a brain in a vat, I would still believe the premise. Similarly, I can't know that the newspaper report is veridical inferentially on the basis of the evidence provided by the Bulls' victory because if the conclusion were false—if the report were not veridical—I would still believe the evidence. The inference from A to *I don't falsely believe A* provides an extreme example of this situation: if the conclusion were false—if I falsely believed A, I would, of necessity, still believe the premise, since it's not possible to falsely believe A without believing A. This seems to me to be a powerful intuitive reason for rejecting this inference as a legitimate source of knowledge.

Propositions to the effect that a proposition is not falsely believed do not allow us to concentrate exclusively on knowledge of the truth values of beliefs, since the proposition ~(Bel(A) & ~A) is true not only when A is truly believed, but also when A is not believed. A similar problem afflicts propositions of the form Bel(A) & A, which are false, not only when A is falsely believed, but also when it is not believed.

We face in effect the same situation as in the case of reflective knowledge, and I propose to adopt the same strategy. This involves concentrating on propositions that ascribe a predicate (…*is true*) to an object identified with a definite description (…*is a belief of mine with A as its content*), and assuming that in its negation the description has wide scope. Thus if V stands now for …*is true* and B stands for …*is a belief of mine with A as its content*, the proposition will be symbolised as V ℩x Bx. The idea of knowing V ℩x Bx inferentially on the basis of evidence provided by A is as unattractive as the idea of knowing in this way that I don't falsely believe A, but since it is in general possible to know A and to know that A entails V ℩x Bx, the inference will be able to produce knowledge, according to Transmission*, unless it violates PI. But we can see easily that it does. To see this, we need to compare p(Bel(A) | V ℩x Bx) with p(Bel(p) | ~V ℩x Bx). Satisfying PI would require the former to be greater than the latter, but this is impossible, since p(Bel(A) | ~V ʅx Bx) = 1: I cannot falsely believe that A without believing that A. Notice also that PI doesn't just manage to rule out the case that we want to rule out. It also reflects the source of our intuitive reluctance to accept that knowledge of the truth value of my belief in A can result from an inference from A. The problem is that I wouldn't be less likely to believe the premise if the conclusion were false than if it were true.

7. Bootstrapping

I want to turn now to another form of inference that has received considerable attention in this connection. In an example that Jonathan Vogel borrows from Michael Williams, Roxanne forms the belief that the petrol tank in her car is full (FULL) whenever she sees that the gauge on the dashboard reads F (GAUGE) (Vogel 2000). Her gauge is highly reliable. Hence, if reliable formation or truth tracking is sufficient for knowledge, then the true beliefs that Roxanne forms in this way will have to be accorded the status of knowledge, even though she has no evidence of the reliability of the gauge. Now suppose that when Roxanne sees that the gauge reads F, in addition to coming to believe FULL, she comes to believe GAUGE, i.e. that the gauge reads F. Let's assume that Roxanne can come to know GAUGE in this way.

What would an inductive argument for the reliability of the gauge look like? Notice first that the claim that the gauge is reliable can be understood as the claim that the gauge reading F provides adequate evidence for the hypothesis that the tank is full. Suppose that we cash out this notion in terms of incremental confirmation, construed, as I think it should be (Zalabardo 2009), as a lower bound on the *likelihood ratio*: p(E | H) / p(E | ~H), written LR(H, E). Then the claim that the gauge is reliable is the claim that LR(FULL, GAUGE) is sufficiently high. An inductive argument for this conclusion would seek to derive it from premises concerning observed frequencies. Thus, from the premise that the proportion of F readings to be found among the observed cases in which the tank is full is considerably higher than the proportion of F readings to be found among the observed cases in which the tank is not full, the argument would conclude that LR(FULL, GAUGE) is high.

Let's assume that the premises of this argument provide adequate support for their conclusion—that evidence concerning observed frequencies provides, in suitable circumstances, adequate support for conclusions about probabilities—and let's assume that Roxanne knows this. Clearly, Roxanne could also know the premises of the argument. Her true beliefs as to whether or not the gauge reads F can have the status of knowledge, and since the gauge is reliable, the true beliefs about the contents of the tank that she forms with the help of the gauge will have to be treated as knowledge by reliabilist and truth-tracking accounts. It follows that Roxanne will have knowledge of the observed relative frequencies that figure in the premises of the argument.

Consider now a non-deductive version of Transmission: if (a) S knows E, (b) S knows that E provides adequate non-deductive support for H and (c) S believes H, then S knows H inferentially on the basis of the information provided by E. If we accepted this principle, we would have to conclude that her inference enables Roxanne to know that the gauge is reliable.

This is a highly counterintuitive outcome. Roxanne cannot use this inference to gain knowledge of the reliability of the gauge. The problem with her procedure doesn't concern the argument itself, or the epistemic status of Roxanne's belief in the premises or in the connection between premises and conclusion. Hence accommodating our intuitive rejection of this form of knowledge acquisition would require arguing that even though Roxanne knows the premises of the argument and she knows that the premises support the conclusion, there is another condition on inferential knowledge that she fails to meet.

Much of the recent literature on this topic assumes that we could avoid counting Roxanne as coming to know with her inductive argument that the gauge is reliable only if we invoked a principle that treats knowledge of the reliability of the gauge as a precondition for obtaining from the gauge knowledge of the contents of the tank (Cohen 2002; Van Cleve 2003). But principles along these lines have been accused of leading directly to scepticism and of being incompatible with reliabilist theories of knowledge. I have argued elsewhere that the problems faced by these principles are not as serious as they might seem at first (Zalabardo 2005). But here I want to present a different strategy for ruling out Roxanne's inference as a case of inferential knowledge.

It seems to me that the most intuitive explanation of the inadequacy of Roxanne's inference focuses on the fact that the gauge is the only method at her disposal for ascertaining whether the tank is full. This circumstance should pose no obstacle to her beliefs about the contents of the tank having the status of knowledge, but it should rule out these beliefs as premises in an inference for the reliability of the gauge. I want to argue that the reason why this feature of Roxanne's situation poses a problem is that it severs the connection between Roxanne's belief in the premises and the truth value of the conclusion. The problem is, once more, that PI is not satisfied: Roxanne is no less likely to believe the premises of the argument if the conclusion is false than if it's true. Given Roxanne's state of information, the probability that she will believe the premises of her inductive argument for the reliability of the gauge is not affected by the value of LR(FULL, GAUGE). She will be just as likely to believe that the observed cases in which the gauge reads F are precisely the cases in which the tank is full if these values are low as if they are high.

On Vogel's construal, Roxanne's inductive argument for the reliability of the gauge involves, for every time t at which she forms belief in GAUGE and FULL in the way described, a lemma to the effect that the gauge is reading accurately on that occasion (ACCURATE). Vogel suggests that the epistemic status of Roxanne's belief in ACCURATE is already problematic, but since she has validly inferred ACCURATE from GAUGE and

FULL, he thinks that any shortcoming of the epistemic status of her belief in ACCURATE would also have to affect her belief in FULL (assuming that her knowledge of GAUGE is above suspicion). I share Vogel's misgivings about the epistemic status of Roxanne's belief in ACCURATE. I don't think she can know this proposition inferentially on the basis of the evidence provided by GAUGE and FULL. But my proposal has the resources for securing this result without withholding from Roxanne's belief in FULL the status of knowledge.

We can use the ideas that we presented in our discussion of reflective knowledge to explain why this inference should be ruled out as a case of inferential knowledge. Notice that ACCURATE can be construed as an instance of the veridicality propositions that we considered there. If R denotes a plausible description of the reading, and V is the predicate that ascribes accuracy to it, ACCURATE can be analysed as V ℩x Rx. Hence, in order to determine whether Roxanne's inference gives her inferential knowledge of ACCURATE, we need to compare p(Bel(GAUGE) & Bel(FULL) | V ℩x Rx) with p(Bel(GAUGE) & Bel(FULL) | ~V ℩x Rx). As I argued in section 5, the description in ~V ℩x Rx should be understood as having wide scope. Hence the question that we need to ask is whether Roxanne is less likely to believe GAUGE and FULL if the gauge is reading inaccurately (i.e. GAUGE & ~FULL) than if it is reading accurately (i.e. GAUGE & FULL). And this question should be answered in the negative. So long as GAUGE is true, the probability that Roxanne believes GAUGE and FULL will be high, and unaffected by the truth value of FULL.

Contrast this situation with one in which Roxanne can ascertain whether or not the tank is full independently of the gauge, say, using a dipstick. Intuitively this would make all the difference to Roxanne's ability to gain knowledge of the reliability of the gauge from the argument. And PI registers this difference. In this new scenario, a low value for LR(FULL, GAUGE) will decrease the probability that she believes that the gauge reads F more often in the cases in which the tank is full than in those in which it isn't full. If LR(FULL, GAUGE) is low, i.e. if p(GAUGE | FULL) is not much higher than p(GAUGE | ~FULL), it is likely that there will be nearly as many observed cases in which the gauge reads F among the cases in which the tank is not full as among those in which it is full, and Roxanne, with her dipstick, will be able to detect this. These points about her inference for the reliability of the gauge can also be applied to her inference for the proposition that the gauge is reading accurately on a given occasion. Armed with her dipstick, Roxanne will be more likely to believe GAUGE and FULL if the gauge is reading accurately than if it isn't.[20]

8. Roush on inferential knowledge

In her recent book (Roush 2005), Sherrilyn Roush has defended an account of knowledge for which she uses the label *recursive tracking*. She presents the aspect of her position on which I want to concentrate in the following passage, where she uses the term *Nozick-knows* to refer to the knowledge that results from truth tracking:

> […] on the new view Nozick-knowing is not the only way to know. From what we Nozick-know we can get by known implication to other beliefs that are also knowledge. Thus, to analyze the concept of knowledge I combine the notion of Nozick-knowing with a recursion clause: For subject S and proposition p, S *knows* that p if and only if:
>
> > S Nozick-knows that p
>
> or

[20] PI will not rule out some inferences that appear illegitimate. Consider, e.g., the inference from HANDS to the proposition that I believe I have hands (Bel(HANDS)) and I'm not a brain in a vat. Suppose that the correlation between HANDS and Bel(HANDS) if I'm not a brain in a vat is such that HANDS provides adequate support for Bel(HANDS) & ~BIV. The inference will satisfy PI so long as p(Bel(HANDS) | Bel(HANDS) & ~BIV) is greater than p(Bel(HANDS) | ~(Bel(HANDS) & ~BIV)), but p(Bel(HANDS) | Bel(HANDS) & ~BIV) equals one, and p(Bel(HANDS) | ~(Bel(HANDS) & ~BIV)) will be less than that. Notice that if we decided to treat this inference as legitimate we wouldn't be forced to accept that I can know ~BIV inferentially. I can't know ~BIV inferentially on the basis of the evidence provided by Bel(HANDS) & ~BIV, since the inference doesn't satisfy PI. I am no less likely to believe this evidence if I am a brain in a vat than if I'm not. See, in this connection, Nozick's discussion of the possibility of knowing a conjunction without knowing each of its conjuncts (Nozick 1981: 228). If we weren't prepared to take this route, we would have to settle for treating PI as a partial diagnosis of the family of difficulties that we have discussed.

> p is true, S believes p, and there is a q not equivalent to p such that q implies p, S knows that q implies p, and S knows that q.
>
> According to this analysis, anything that you derive from something you Nozick-know by n steps of deduction, for some finite n, is also something you know. (Roush 2005: 42-43)

This is only a preliminary formulation of Roush's highly sophisticated proposal, but it adequately highlights the features on which I want to focus.

There are some aspects of this position that I find very appealing. I agree with Roush that truth tracking should be treated as a sufficient condition for knowledge, but we should be able to have inferential knowledge in cases in which we don't track the truth. There are, however, two important respects in which my views differ from hers.

The first is that her recursion clause contemplates inferential knowledge involving only *deductive* inference. Notice that this feature of Roush's account doesn't entail that knowledge cannot be acquired by non-deductive inference. What it does entail is that this will be possible only when as a result of the inferential process the subject's belief in the conclusion comes to track the truth. Roush is explicit about this: "it is clear that on this view all inductive routes to knowledge must be such that through them we satisfy the tracking conditions" (Roush 2005: 52). I find this aspect of Roush's position unsatisfactory. It seems to me that non-deductive inference should also enable us to obtain inferential knowledge propositions whose truth we don't track, but I am not going to defend this point here.

The second aspect of Roush's position that I find unsatisfactory is directly connected to the issues that I have discussed in this paper. Roush is prepared to accept as knowledge all cases in which I know that a proposition I believe is deductively entailed by known evidence, whereas I have argued that we shouldn't treat in this way cases in which PI is not satisfied. As we have seen, the main consequence of this restriction is to rule out three types of case: Moorean inferences, inferences from a belief to the veridicality of its source (or to the truth of the belief), and inductive bootstrapping arguments. Roush discusses all these cases in some detail.

Concerning Moorean inferences, there can be no question that it follows from Roush's recursive tracking account of knowledge that if my belief in HANDS tracks the truth, and I know that HANDS entails ~BIV, then I know ~BIV. Furthermore, on her account, knowledge of ~BIV would be "gained via known implication from beliefs that are already knowledge" (Roush 2005: 51).

Nevertheless, Roush seems reluctant to accept this consequence of her view. She writes:

> According to this view of knowledge I may know that there is a table in front of me, in which case I also know that I am not a brain in a vat (by known implication), or I may not know that there is a table in front of me, because I do not know that I am not a brain in a vat. Recursive tracking does not determine which of these positions one must adopt […] (Roush 2005: 55).

Her choice of everyday proposition in this passage is unfortunate, since the proposition that there is a table in front of me does not entail ~BIV, but let's assume for the sake of the argument that the entailment holds. My main point about this passage is that if I believe that my belief that there is a table in front of me tracks the truth, then, contrary to what Roush suggests, recursive tracking does tell me which of these propositions to adopt: I have to believe that I also know ~BIV (by known implication). If Closure is treated as an independent constraint on our knowledge ascriptions, then it is indeed neutral as between the two options that Roush describes. But the same cannot be claimed for recursive tracking. If I believe that my belief in HANDS tracks the truth, and that I know that HANDS entails ~BIV, then recursive tracking doesn't leave me the option of saying that I don't know HANDS because I don't know ~BIV. It forces me to say that I know HANDS (by tracking) and ~BIV (by known implication).

Roush's discussion of Moorean inferences reveals another important consequence of her account. It seems natural to suppose that, when a proposition p is known inferentially, the evidence on which this knowledge is based will provide the subject with adequate reasons or justification for p. Roush, however, doesn't expect her inductive clause to throw any light in general on the justificatory status of beliefs. She writes:

> […] though we may, and ordinarily think we do, have knowledge that we are not brains in vats, we lack, and will always lack, the ability to offer justification for such claims. This would mean that there can be knowledge without justification, a view that I hold on other grounds […]. (Roush 2005: 56)

On her view, knowledge of ~BIV is gained by known implication from HANDS, even though HANDS provides no justification for ~BIV. It seems to me that it would be desirable to preserve the link between justification and inferential knowledge, and I want to suggest that PI is a step in this direction, as the inferences that violate it don't seem to provide the subject with justification for their conclusions.

Concerning reflective knowledge, once again, it seems hard to deny that recursive tracking dictates that, if my belief that p tracks the truth, and I know that p entails that my belief that p is true,[21] I will count as knowing that my belief that p is true by known implication. Roush accepts that I can come to know that my belief that p is true in this way for any p that I know, but once again she seems reluctant to accept this consequence of her view. This reluctance is manifested in her discussion of someone's belief that there is no motion of the earth relative to the ether. Assume that this belief is knowledge and that the subject believes that it is not false. Roush writes:

> Still, though it does seem possible that her reflective belief is knowledge we seem to need to know more than that she knows p in order to see her as knowing that she does not falsely believe p, even when she believes the latter. […] I conclude that it ought to follow from a view of knowledge that there are ways of acquiring the reflective knowledge in question but that it is not automatic, and less effort may be needed for this in the case of easily known statements like 'I have hands' than is required for more elaborate beliefs whose status as knowledge itself required much more deliberate

[21] Here and elsewhere I am assuming that I know that I have the belief whose truth value is at issue.

> effort (on the part of someone, not necessarily the subject). (Roush 2005: 60)

I agree that it ought to follow from a view of knowledge that the acquisition of reflective knowledge should not be automatic. However it is clear that Roush's view does not satisfy this requirement. It follows from her view that if I know p, and I believe that my belief that p is true (or not false), the reflective belief will have the status of knowledge so long as I know that p entails the corresponding reflective belief. It seems to me that knowledge of this entailment is a sufficiently weak requirement for the resulting reflective knowledge to count as automatic, since knowledge of the entailment would seem to be required for possessing the concept of true belief, which is required, in turn, for having reflective beliefs. At any rate, knowledge of the entailment won't be harder in cases in which knowledge of p requires deliberate efforts than in cases in which p is easily known, as Roush thinks it should be.

Notice that this route to automatic reflective knowledge is blocked by PI. Knowing p and knowing that p entails that my belief that p is true is not sufficient for inferential knowledge of the reflective proposition, since, as we have seen, p does not provide me with adequate evidence for the proposition that my belief that p is true.

Let's turn now to bootstrapping inferences for reliability claims. As Roush explains, recursive tracking is not committed to the view that these inferences confer the status of knowledge on their conclusions. Roxanne's belief that the petrol gauge in her car is reliable doesn't come to track the truth as a result of her inference, and it doesn't acquire the status of knowledge by known implication either, since as Roush points out, the inference involves non-deductive steps.

Notice, however, that recursive tracking *is* committed to conferring the status of knowledge on Roxanne's beliefs to the effect that the gauge is reading accurately on particular occasions. As Roush puts it,

> The steps of S's procedure that are deductive—conjunction and the inference from 'F and the gauge says "F"' to 'the gauge was accurate this time'—cannot be objectionable to recursive tracking, which allows that knowledge is preserved by deduction. (Roush 2005: 120)

I have argued above that this outcome is in conflict with our intuitions. Roxanne might know these propositions, and knowing them might be required for gaining knowledge from the gauge about the contents of the tank. What I find implausible is the idea that Roxanne can come to know that the gauge is reading accurately by virtue of the fact that she knows this proposition to be entailed by the propositions that the gauge is reading F and the tank is full, which she also knows. PI explains why this account of how Roxanne knows that the gauge is reading accurately cannot be right. Recursive tracking, by contrast, treats it as the right account.

APPENDIX

Theorem: If p(A | B) > p(A | C) and B & C is logically false, then p(A | B) > p(A | C $\vee$ B).

Proof:

p(A | B) > p(A | C)

↓ (by the definition of conditional probability)

$$\frac{p(A \,\&\, B)}{p(B)} > \frac{p(A \,\&\, C)}{p(C)}$$

↓

p(A & B) · p(C) > p(A & C) · p(B)

↓

p(A & B) · p(C) + p(A & B) · p(B) > p(A & C) · p(B) + p(A & B) · p(B)

↓

p(A & B) · (p(C) + p(B)) > p(B) · (p(A & C) + p(A & B))

↓

$$\frac{p(A \,\&\, B)\cdot(p(C)+\ p(B))}{p(B)\cdot(p(C)+\ p(B))} > \frac{p(B)\cdot\big(p(A \,\&\, C)+\ p(A \,\&\, B)\big)}{p(B)\cdot(p(C)+\ p(B))}$$

↓

$$\frac{p(A \,\&\, B)}{p(B)} > \frac{p(A \,\&\, C)+\ p(A \,\&\, B)}{p(C)+\ p(B)}$$

↓ (by the addition axiom, since C & B is logically false)

$$\frac{p(A \,\&\, B)}{p(B)} > \frac{p\big((A \,\&\, C)\vee(A \,\&\, B)\big)}{p(C)+\ p(B)}$$

↓

$$\frac{p(A \,\&\, B)}{p(B)} > \frac{p(A \,\&\, (C\vee B))}{p(C)+\ p(B)}$$

↓ (by the addition axiom, since C & B is logically false)

$$\frac{p(A \,\&\, B)}{p(B)} > \frac{p(A \,\&\, (C\vee B))}{p(C\vee B)}$$

↓ (by the definition of conditional probability)

$p(A \mid B) > p(A \mid B \vee C)$

REFERENCES

Brown, Jessica. 2003. "The Reductio Argument and Transmission of Warrant". In *New Essays on Semantic Externalism and Self-Knowledge*, edited by S. Nuccetelli. Cambridge, Massachusetts: MIT Press.

Brueckner, Anthony. 1994. "The Structure of the Skeptical Argument". *Philosophy and Phenomenological Research* 54:827-35.

Cohen, Stewart. 2002. "Basic Knowledge and the Problem of Easy Knowledge". *Philosophy and Phenomenological Research* 65:309-29.

Davies, Martin. 1998. "Externalism, Architecturalism, and Epistemic Warrant". In *Knowing Our Own Minds*, edited by C. Wright, B. Smith and C. Macdonald. Oxford: Oxford University Press.

DeRose, Keith. 1995. "Solving the Skeptical Problem". *Philosophical Review* 104:1-52.

Dretske, Fred. 1970. "Epistemic Operators". *Journal of Philosophy* 67:1007-23.

Glynn, Luke. 2010. "Deterministic Chance". *British Journal for the Philosophy of Science* 61:51-80.

Goldman, Alvin I. 1986. *Epistemology and Cognition*. Cambridge, Mass.: Harvard University Press.

Hoefer, Carl. 2007. "The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance". *Mind* 116:549-996.

Lewis, David. 1986. "A Subjectivist's Guide to Objective Chance". In *Philosophical Papers, Volume II*. Oxford: Oxford University Press.

———. 1994. "Humean Supervenience Debugged". *Mind* 103:473-90.

Luzzi, Federico. 2010. "Counter-Closure". *Australasian Journal of Philosophy* 88: 673-83.

Moore, G. E. 1939. "Proof of an External World". *Proceedings of the British Academy* 25:273-300.

Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge, Massachusetts: Harvard University Press.

Plantinga, Alvin. 1993. *Warrant: The Current Debate*. New York and Oxford: Oxford University Press.

Pryor, James. 2004. "What's Wrong With Moore's Argument?". *Philosophical Issues* 14:349-78.

Roush, Sherrilyn. 2005. *Tracking Truth*. Oxford: Oxford University Press.

Van Cleve, James. 2003. "Is Knowledge Easy-or Impossible? Externalism as the Only Alternative to Skepticism". In *The Skeptics. Contemporary Essays*, edited by S. Luper. Aldershot, Hampshire: Ashgate.

Vogel, Jonathan. 2000. "Reliabilism Leveled". *Journal of Philosophy* 97:602-23.

Warfield, Ted A. 2005. "Knowledge from Falsehood". *Philosophical Perspectives* 19:405-16.

Wright, Crispin. 1985. "Facts and Certainty". *Proceedings of the British Academy* 71:429-72.

———. 2000. "Cogency and Question-Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof". *Philosophical Issues* 10:140-63.

———. 2002. "(Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell". *Philosophy and Phenomenological Research* 65:330-48.

Zalabardo, José L. 2005. "Externalism, Skepticism and the Problem of Easy Knowledge". *Philosophical Review* 114:33-61.

———. 2009. "An Argument for the Likelihood-Ratio Measure of Confirmation". *Analysis* 69:630-35.

———. 2012. *Scepticism and Reliable Belief*. Oxford: Oxford University Press.

———. forthcoming. "Wright on Moore". In *Wittgenstein, Epistemology and Mind. Themes from the Philosophy of Crispin Wright*, edited by A. Coliva. Oxford: Oxford University Press.