

volume IX ♦ issue 4
2018

Logos & Episteme

an international journal
of epistemology

**Romanian Academy
Iasi Branch**



**“Gheorghe Zane” Institute
for Economic and Social
Research**

Editorial board

Director

Teodor Dima

Editor-in-Chief

Eugen Huzum

Executive Editors

Vasile Pleșca

Cătălina-Daniela Răducu

Assistant Editors

Irina Frasin

Bogdan Ștefanachi

Ioan Alexandru Tofan

Web&Graphics

Codrin Dinu Vasiliu

Virgil-Constantin Fătu

Simona-Roxana Ulman

Contact address:

Institutul de Cercetări
Economice și Sociale „Gh.Zane”

Iași, str.T.Codrescu, nr.2, cod 700481

Tel/Fax: 004 0332 408922

Email: logosandepisteme@yahoo.com

www.logos-and-episteme.projectsbc.ro

Advisory board

Frederick R. Adams

University of Delaware, USA

Scott F. Aikin

Vanderbilt University, USA

Daniel Andler

Université Paris-Sorbonne, Paris IV, France

Alexandru Boboc

Academia Română, Romania

Panayot Butchvarov

University of Iowa, USA

Mircea Dumitru

Universitatea din București, Romania

Sanford Goldberg

Northwestern University, Evanston, USA

Alvin I. Goldman

Rutgers, The State University of New Jersey, USA

Susan Haack

University of Miami, USA

Stephen Hetherington

The University of New South Wales, Sydney, Australia

Paul Humphreys

University of Virginia, USA

Jonathan L. Kvanvig

Baylor University, USA

Thierry Martin

Université de Franche-Comté, Besançon, France

Jürgen Mittelstraß

Universität Konstanz, Germany

Christian Möckel

Humboldt-Universität zu Berlin, Germany

Maryvonne Perrot

Université de Bourgogne, Dijon, France

Olga Maria Pombo-Martins

Universidade de Lisboa, Portugal

Duncan Pritchard

University of Edinburgh, United Kingdom

Nicholas Rescher

University of Pittsburgh, USA

Rahman Shahid

Université Lille 3, France

Ernest Sosa

Rutgers University, USA

Alexandru Surdu

Academia Română, Romania

John F. Symons

University of Texas at El Paso, USA

Logos &
Episteme

TABLE OF CONTENTS

RESEARCH ARTICLES

Umut BAYSAN, Memory, Confabulation, and Epistemic Failure.....	369
John BIRO and Fabio LAMPERT, 'Peer Disagreement' and Evidence of Evidence.....	379
Miguel LÓPEZ-ASTORGA, An Axiom Linking Necessity and Obligation Provided by Prior and Its Analysis Under Carnap's Method.....	403
P.D. MAGNUS, Science, Values, and the Priority of Evidence.....	413
Timothy PERRINE, Evidentialism, Knowledge, and Evidence Possession.....	433
Mattias SKIPPER, Higher-Order Defeat Without Epistemic Dilemmas.....	451
Jacques-Henri VOLLET, The Warrant Account and the Prominence of 'Know'.....	467
Notes on the Contributors.....	485
<i>Logos and Episteme</i> . Aims and Scope.....	487
Notes to Contributors.....	489

RESEARCH ARTICLES

MEMORY, CONFABULATION, AND EPISTEMIC FAILURE

Umut BAYSAN

ABSTRACT: Mnemonic confabulation is an epistemic failure that involves memory error. In this paper, I examine an account of mnemonic confabulation offered by Sarah Robins in a number of works. In Robins' framework, mnemonic cognitive states in general (e.g., remembering, misremembering) are individuated by three conditions: existence of the target event, matching of the representation and the target event, and an appropriate causal connection between the target event and its representation. Robins argues that when these three conditions are not met, the cognitive state in question is an instance of mnemonic confabulation. Here, I argue that this is not true. There are mnemonic cognitive states which don't meet any of these conditions, and they are not cases of mnemonic confabulation. On a more positive note, I argue that mnemonic confabulation requires it to be a failing on behalf of either the subject or her mnemonic system that these conditions are not met.

KEYWORDS: confabulation, epistemic failure, memory, misremembering

1. Introduction

Confabulation is an epistemic failure, and in paradigmatic cases, it involves failure of remembering. In what ways a subject's remembering has to fail in order for her to count as confabulating a memory is a venue for philosophical debate. This paper aims to contribute to this debate. In what follows, I examine an account of confabulation proposed by Sarah Robins,¹ and argue that although the background philosophical framework that Robins has developed is commendable, her treatment of confabulation yields some counterintuitive results. In her work, Robins focuses on *mnemonic confabulation* (instead of confabulation *simpliciter*), and here I will follow her in doing that (except when I explicitly state otherwise).

¹ Sarah Robins, "Misremembering," *Philosophical Psychology* 29 (2016): 432-447, "Confabulation and Constructive Memory," *Synthese* (2017): 1-17, and "Mnemonic Confabulation" (unpublished manuscript).

2. Robins on Mnemonic Cognitive States

Mnemonic confabulation is the type of confabulation that involves a memory error. To understand Robins' treatment of mnemonic confabulation, let's explore the ways she contrasts mnemonic confabulation with other mnemonic cognitive states.

The paradigm case of a mnemonic cognitive state is successfully remembering a past episode. (Taking remembering to be a factive mental state, I will henceforth omit "successfully."²) Suppose Jude remembers that Sue bought him cufflinks for his 31st birthday. Let's say that Jude very vividly remembers the occasion with its relevant details: its being his birthday, and in fact being his 31st birthday, that Sue gave him a navy blue box just when the main dish was being served, that he opened the box and there were a pair of cufflinks, and so on. According to Robins, this instance of remembering involves, *first*, the fact that the target event did take place (i.e., Jude had a 31st birthday which he celebrated with his partner, and Sue gave him cufflinks on that very occasion); *second*, that Jude has a mental representation of the target event in a way that matches the target event with respect to its relevant aspects (i.e., the event is represented as a birthday celebration event, as well as a Sue-gifting-Jude-cufflinks event); and *third*, the right kind of causal connection between the target event and the mental representation of the target event (i.e., Jude's birthday celebration and the cufflinks-giving event are among the causal antecedents of the representation). Let's call these three conditions TARGET, MATCHING, and CAUSATION respectively.

The first two conditions (TARGET and MATCHING) can account for the fact that remembering is a factive mental state: the mental state that underlies remembering correctly represents the remembered event which indeed took place in the way it is represented. The third condition (CAUSATION) rules out cases where the representation and the target event match, but not for the right reasons. Suppose it is true that Sue bought Jude cufflinks for his 31st birthday, but Jude never came to know this. But a demon implants a chip in his brain which physically realizes a mnemonic mental representation that has the relevant aspects of the target event. Although he has an accurate representation of a past event, this

² I assume, without argument, that remembering is factive. I acknowledge that this could be debated.

shouldn't count as remembering. CAUSATION rules out this case from being a case of remembering.

Another type of mnemonic cognitive state is misremembering. Imagine that Jude seems to remember that Sue bought him cufflinks with blue prints on them, whereas in fact the prints were burgundy. Apart from this difference, let's say that Jude's mental representation of his 31st birthday is accurate. Although, in this case, he still remembers certain aspects of the target event, overall, his cognitive state counts as misremembering. In this case, his representation still picks out the target event (i.e., his 31st birthday celebration), but it misrepresents some of its aspects. In cases of misremembering, despite the satisfaction of TARGET (i.e., the target event exists), MATCHING is not satisfied (i.e., the content and the target event do not match). Thus, misremembering is not factive. The difference between remembering and misremembering is similar to the difference between veridical perceptions and perceptual illusions.³ In both veridical perceptions and perceptual illusions, a target object does exist, but whereas in the former, the target object is perceived as it is (i.e., it does have the sensory properties it appears to have), in the latter, the target object is misperceived (i.e., it doesn't have some of the sensory properties it appears to have).

Now, we are in a position to understand Robins' account of mnemonic confabulation. Mnemonic confabulation is also a mnemonic cognitive state, and it is different from both remembering and misremembering in important aspects. Suppose that Sue did not buy Jude anything for his 31st birthday, and in fact, they did not even have any celebration. Now suppose Jude seems to have a memory of Sue giving him cufflinks on his 31st birthday. So, he has a mental representation of an event which has the aspects of a 31st birthday event, a celebration dinner event, a cufflinks-gifting event, and so on. According to Robins, this would be an example of mnemonic confabulation. This state is importantly different from both remembering and misremembering. Unlike in cases of both remembering and misremembering, the target event does not exist. There is no 31st birthday celebration regardless of whether cufflinks were given or not. So, TARGET is not satisfied. Given that the target event does not exist, MATCHING is not satisfied either: the target event and the content of the representation do not match (simply because the target event doesn't exist). And furthermore, since these two conditions are not met, CAUSATION fails also: there is no right kind of causal relation between a target event and the content of the representation. Thus,

³ Robins, "Confabulation and Constructive Memory."

mnemonic confabulation differs from remembering and misremembering in the sense that all three conditions for mnemonic cognitive states fail to be satisfied (whereas in remembering all three are satisfied, and in misremembering, at least the first one is satisfied). Here, Robins compares mnemonic confabulation to cases of perceptual hallucinations. In perceptual hallucinations, there is no target object although there is a representation of an object with some sensory properties. Likewise, in mnemonic confabulation, there is no target event despite the fact that there is a representation of an event with certain aspects.

What we see here is an elegant framework which locates different sorts of mnemonic cognitive state in one table. Mnemonic states are representational states, and in individuating different kinds of mnemonic cognitive states, the relevant parameters are TARGET, MATCHING, and CAUSATION. The account is also in line with a more general framework according to which memory and other mnemonic states are understood in terms of causation.⁴ The following falls out from Robins' account as a characterisation of mnemonic confabulation:

(MC) A subject mnemonically confabulates some putative past event if and only if she has a mnemonic representation which meets none of TARGET, MATCHING, and CAUSATION.⁵

Although this is an elegant framework and arguably successful in explaining contrasting remembering and misremembering, I believe that **(MC)** yields counterintuitive results, which I shall highlight next.

3. Some Counterintuitive Results

I believe that a central aspect of confabulation is missing in **(MC)**. I will illustrate this by giving an example which would count as mnemonic confabulation according to **(MC)**, and then argue that it shouldn't. That will show that being a mnemonic mental representation that satisfies none of TARGET, MATCHING, and CAUSATION is *not sufficient* for being an instance of mnemonic confabulation.

⁴ See also Sven Bernecker, "A Causal Theory of Mnemonic Confabulation," *Frontiers in Psychology* 8 (2017): 1205.

⁵ Or when it does meet TARGET, it is purely accidental that it does. Bernecker's (*ibid.*) account of confabulation emphasises this possibility. He suggests that "a piece of confabulation may even be entirely correct. It is possible that a patient fantasizes correctly by telling a story that, by sheer luck, represents the objective reality" (*ibid.*, 5). In fact, Bernecker uses this and similar considerations to argue that the real mark of mnemonic confabulation is the failure of what I have in this paper called CAUSATION.

After considering and responding to possible ways this objection could be replied, I will remark on what I think is missing from this account.

Consider the following (very dull) story which I shall call *the flapjack case*. At t_1 , I am sitting in a café, sipping my coffee, and it appears to me that there is a piece of flapjack on a plate on the table opposite of me. Actually, there is no flapjack on the table, and in fact there is nothing on the table. So, I am hallucinating a piece of flapjack on the opposite table. (It doesn't matter what causes this hallucination; to fix ideas, let's suppose it is a malicious demon behind this very uninteresting trick.) There is nothing suspicious about there being some flapjack on a table; I am in a café after all, and many cafes do serve flapjack. So, I have no reason to doubt the veridicality of this experience or reflect much on it. Days pass, now the time is t_2 . I am sitting in another café, some stranger approaches and offers me a flapjack. This prompts me to recall an experience I had recently. Then I form the mental representation of a flapjack-on-the-opposite-table event that happened at t_1 .

According to (MC), the mnemonic cognitive state I am in at t_2 should count as mnemonic confabulation. The target event doesn't exist; there wasn't a flapjack-on-the-opposite-table event. Since there was no such event at t_1 , the content of the mental representation at t_2 doesn't match a target event at t_1 ; and for the same reason, there is no right kind of causal connection between an event at t_1 and the representation at t_2 . So, there is a mental representation that doesn't meet any of TARGET, MATCHING, and CAUSATION. However, intuitively, it is not right that this case is a case of mnemonic confabulation. Therefore, (MC) doesn't capture the essence of mnemonic confabulation.

One might think that the target event does exist in the flapjack case; it is just not a flapjack-on-the-opposite-table event. As a sitting-in-café event at t_1 , the target event does exist, it might be argued. Whether this response is viable partly depends on how to individuate events. If the target event is essentially a café event, then the target event indeed exists. But in the representation of this event at t_2 , the salient feature of it is that it is a flapjack-on-the-opposite-table event, which suggests that it is more appropriate to take it as an essentially flapjack-on-the-opposite-table event. Nevertheless, even if the target event is essentially a sitting-in-café event and hence that the target event at t_1 does exist, Robins' account runs into a different problem. For the sake of entertaining this response, let's accept that the target event does exist at t_1 as a sitting-in-café event, but it is misrepresented at t_2 as a flapjack-on-the-opposite-table event. From the characterisation of

mnemonic cognitive states Robins gives, this state then should be categorised as misremembering: TARGET is satisfied, MATCHING is not. But this is an equally implausible consequence. Misremembering should be a failure of remembering. Here, there is nothing that indicates that the failure has anything to do with remembering. The event at t_1 has always been a flapjack event for me.

4. Further Possible Responses

We have seen an example of a mnemonic cognitive state which fails all three conditions TARGET, MATCHING, and CAUSATION, yet it makes little sense to categorise it as a case of mnemonic confabulation. Before remarking on what I think is missing from this example to make it a case of confabulation, let me address two possible responses that Robins can offer.

First, Robins can argue that, in the flapjack case, the target event does indeed take place. This is not because there is a sitting-in-café event (*à la* the misremembering response discussed above), but it is due to the fact that the target event is a sensory experience. That is, it is true that there is no external object (i.e., flapjack) at t_1 , but there is a sensory object that exists at t_1 . So, Robins can argue that the flapjack case doesn't count as mnemonic confabulation according to (MC) because TARGET is satisfied. Let's call this *the sensory event response*.

I don't think that the sensory event response is a satisfactory one. Why posit internal sensory objects just to get around this particular type of counterexample? If we are to posit an internal sensory object to explain the flapjack case, why not take the target event in cases of remembering also as internal sensory events? If we are to hold that when I remember that I saw a deer in the forest what I remember is not a deer but instead an internal sensory object (a deer-like sense-datum), why not also hold that when I take myself to see a deer in the forest, what I in fact perceive is a deer sense-datum? This is not the place to give an argument against the sense-data theory, but it is worth noting that this response comes with the burden of making a case for sense-data.

Regardless of any qualms about the metaphysics of sensory objects, it is not clear that it is a good move to suggest that target events are sensory events in *all* mnemonic cognitive states (which we should do if we want to make the sensory event response sound less *ad hoc*). We would then be holding that, in a case of remembering, there is an external event, which then causes a sensory event, which in turn is accurately represented (and representation is caused in the right way by the sensory event). What is problematic with this is that the representation would

be a *representation of a sensory event, not a representation of an external event*. Introspectively, I find this hard to believe. On a more theoretical point, it is problematic to think that at every time I seem to remember a physical event e , I actually remember a sensory event e^* , but I take myself to remember e (or I am in a position to take myself to remember e). If remembering, as a mental state kind, is to accommodate this possibility, it shouldn't be categorised as a factive mental state kind. There surely are cases when we recall our sensory experiences; but when we do so, we remember them *as sensory experiences*. If we don't remember them as sensory experiences, we don't remember them *simpliciter*. I believe these difficulties make it very difficult to motivate the sensory event response.

So much for the sensory event response. What about Robins' second possible move? Robins can bite the bullet and hold that the flapjack case is indeed a case of mnemonic confabulation. However, note how different this case would be from more paradigmatic cases of mnemonic confabulation. In the paradigmatic cases of mnemonic confabulation, the inaccurate representation is a failing on part of the mnemonic system. Mnemonic confabulation is a failure of remembering. In the flapjack case, there is no failure of remembering. If there is any failure, it has to do with the forming of the experience at t_1 in the first place.

I believe what we have just seen reveals what is missing in the characterisation of mnemonic confabulation in (MC). It is evident (from the fact that she compares mnemonic confabulation to perceptual hallucinations) that Robins (rightly) treats mnemonic confabulation as an epistemically unideal kind of mental state. Mnemonic confabulation involves a form of failure. However, (MC) doesn't give us any clue as to where that failure lies.

5. Normativity in Confabulation

I mentioned that I am following Robins in focusing on mnemonic confabulation rather than confabulation *simpliciter*. In this section and next, I shall relate the discussion so far to the concept of confabulation, broadly understood. My intuitions regarding the flapjack case are motivated by the fact that I take mnemonic confabulation to be a specific kind of confabulation. I think a theory of mnemonic confabulation would be unattractive if it couldn't accommodate the fact that mnemonic confabulation is a type confabulation. But what is confabulation more generally?

The nature of confabulation is a matter of dispute among philosophers of cognitive science, philosophers of psychiatry, psychiatrists, and others.⁶ Whereas some researchers restrict the term “confabulation” to epistemically problematic mental states which have to do with memory,⁷ others take confabulation to be a more general epistemic failure which involves false beliefs regardless of whether these false beliefs concern putative past events or not.⁸ But surely, not every false belief counts as confabulation.⁹ If it did, then why need the category of confabulation over and above the category of false belief? Then, what is the additional component in confabulation on top of a false belief?

Researchers seem to agree that one of the things that mark the difference between a merely false belief and a confabulatory mental state is that in the latter, there is failing on behalf of either the subject or the subject’s mnemonic system where there ideally shouldn’t be. Turnbull and colleagues suggest that, in confabulation, “false beliefs and opinions about the world ... are manifestly incorrect.”¹⁰ Being manifestly incorrect, these beliefs are beliefs that the subject should not have formed or retained. Hirstein suggests that when a subject confabulates that P, her belief that P is ill-grounded and moreover that subject should (but does not) know that her belief is ill-grounded.¹¹ These suggest that confabulation, if it is to be separated from a merely false belief, involves a normative element.

What do I mean by a normative element? When I say that a confabulated belief is a belief that should not have been formed or retained, am I suggesting that the subject had an obligation not to form that belief? If the idea of obligation is

⁶ For discussion, see Lisa Bortolotti’s *Delusions and other Irrational Beliefs* (Oxford: Oxford University Press, 2010), 43-50.

⁷ For example, Aikaterini Fotopoulou, “False-Selves in Neuropsychological Rehabilitation: The Challenge of Confabulation,” *Neuropsychological Rehabilitation* 18 (2008): 541-565.

⁸ Oliver H. Turnbull, Sarah Jenkins, and Martina L. Rowley, “The Pleasantness of False Beliefs: An Emotion-Based Account of Confabulation,” *Neuro-Psychoanalysis* 6 (2004): 5-45, William Hirstein, *Brain Fiction: Self-deception and the Riddle of Confabulation* (Cambridge, MA: MIT Press, 2005), Linda Örluv and Lars-Christer Hydén, “Confabulation: Sense-Making, Self-Making and World-Making in Dementia,” *Discourse Studies* 8 (2006): 647-673, and G.E. Berrios, “Confabulations”, in *Memory Disorders in Psychiatric Practice*, eds. G.E. Berrios and J.R. Hodges (New York, NT: Cambridge University Press, 2000), 348-368.

⁹ Also, as noted in footnote 5, it is possible for the content of a confabulation to be accidentally true. I may confabulate that P whereas P happens to be true (as in Gettier cases).

¹⁰ Turnbull, Jenkins, and Rowley, “The Pleasantness of False Beliefs,” 6.

¹¹ Hirstein, *Brain Fiction*, 187.

linked to that of responsibility, does this mean that, when S confabulates a belief, S *could have* believed otherwise? Ideally, I would like not to make any of these commitments. After all, it is plausible that, at least in some cases of confabulation in the clinical population, subjects could not have done otherwise. This suggests that the sense of normativity here is different from the sense of normativity that underlies moral responsibility. Nevertheless, it is clear that, in cases of confabulation (likewise in cases of delusions and irrational beliefs), there is a sense in which either subjects or their mnemonic systems depart from some epistemic norms.

Could CAUSATION in Robins' account not be viewed as a normative requirement? In cases of remembering, representations must be *appropriately* caused. It might be thought that the appropriateness of the causal connection could underwrite the normative element that I argue is missing in Robins' account. The problem with this suggestion is that, in (MC), CAUSATION fails purely in virtue of the failure of TARGET. Its failure has nothing to do with the causal connection appropriateness. So, overall, I don't think that (MC) captures the required normativity.

6. Concluding Remarks and a Proposal

It strikes me as evident that (MC) fails to have the normative element required from an account confabulation. In concluding, let me highlight four possible ways this may have bearing on Robins' account. First, and least desirably, one could just argue that because mnemonic confabulation, as discussed above, fails to be a form of confabulation due to failing to meet a normativity criterion, Robins' framework should be abandoned altogether. Although I am mentioning this possibility, let me make it explicit that this is not the recommendation I am making; there are less radical ways to resolve the issue at hand. Second, and less undesirably, one could just accept that mnemonic confabulation is a very different type of cognitive state compared to confabulation simpliciter. At this point, it may be merely a terminological dispute as to whether mnemonic confabulation should be called as such. Third, and relatedly, one could argue that mnemonic confabulation is a type of confabulation, but the class of confabulatory mental states are very diverse. If this is the preferred option, one should also be prepared to respond to some worries with respect to whether confabulation, as a mental state kind, is a natural kind or not. And finally, one could agree with the message of the previous section, and accept that confabulation has to have a normative element. If mnemonic

confabulation is a type of confabulation, then mnemonic confabulation must have a normative element too. If that is the case, the spirit of Robins' account can be retained, but can be supplemented with a normative criterion. One way of doing so would be to hold the following:

(MCN) A subject mnemonically confabulates some putative past event if and only if she has a mnemonic representation which meets none of TARGET, MATCHING, and CAUSATION, and it is a failing on behalf of either the subject or her mnemonic system that none of these conditions is met.¹²

This might not be the only way the problem I have highlighted could be solved, but it is one way of solving it, and I hope it is helpful way of doing so.¹³

¹² **(MCN)** is an account of mnemonic confabulation, but admittedly it fails to accommodate the possibility of veridical mnemonic confabulation, a possibility that one might want to consider as per footnotes 5 and 9 above. To get around this problem, we can add a disjunctive clause in **(MCN)** to the effect that when TARGET and MATCHING are met, it is only accidental (as in Gettier cases) that they are met.

¹³ Thanks to Kathy Puddifoot for her helpful comments on a previous version of this paper.

‘PEER DISAGREEMENT’ AND EVIDENCE OF EVIDENCE

John BIRO, Fabio LAMPERT

ABSTRACT: What the rational thing to do in the face of disagreement by an epistemic peer is has been much discussed recently. Those who think that a peer’s disagreement is itself evidence against one’s belief, as many do, are committed to a special form of epistemic dependence. If such disagreement is really evidence, it seems reasonable to take it into account and to adjust one’s belief accordingly. But then it seems that the belief one ends up with depends, in part, on what someone else believes, even if one does not know why that someone believes what he does. While the practical impossibility of finding actual cases of peer disagreement has been often noted, its conceptual possibility has gone unquestioned. Here we challenge this consensus and argue, first, that, strictly speaking, peer disagreement is impossible and, second, that cases of – all-too-common – near-peer disagreement present no special puzzle and require nothing more than adhering to standard principles of sensible epistemic conduct. In particular, we argue that in such cases there is no good reason to adopt the widely accepted principle that evidence of evidence is evidence. If so, even if one takes a near-peer’s disagreement as a reason for re-examining one’s belief, one is not epistemically dependent in the sense one would be if that disagreement were evidence concerning the matter in question.

KEYWORDS: disagreement, evidence, epistemic peers, justification, rationality

Introduction

What the rational thing to do in the face of disagreement by an epistemic peer is has been much discussed recently. Those who think that a peer’s disagreement is itself evidence against one’s belief, as many do, are committed to a special form of epistemic dependence. If such disagreement is really evidence, it seems reasonable to take it into account and to adjust one’s belief accordingly. But then it seems that the belief one ends up with depends, in part, on what someone else believes, even if one does not know why that someone believes what he does. While the practical impossibility of finding actual cases of peer disagreement has been often noted, its conceptual possibility has gone unquestioned. Here we challenge this consensus and argue, first, that, strictly speaking, peer disagreement is impossible and, second, that cases of – all-too-common – near-peer disagreement present no special

puzzle and require nothing more than adhering to standard principles of sensible epistemic conduct. In particular, we argue that in such cases there is no good reason to adopt the widely accepted principle that evidence of evidence is evidence. If so, even if one takes a near-peer's disagreement as a reason for re-examining one's belief, one is not epistemically dependent in the sense one would be if that disagreement were evidence concerning the matter in question.

Disagreement is often taken to be puzzling if it involves these elements: two parties with equal intellectual virtues and abilities, the same evidence regarding some proposition p , full disclosure of the evidence between the parties, yet disagreement over the truth-value of p . Here are some well-known examples:

Feldman's Quad: Two people are standing by the window looking out on the quad. They seem to be equal in intellectual virtue and cognitive ability. One claims to be seeing the Dean standing in the middle of the quad, the other disagrees. They are both confident in their beliefs, taking themselves to be correct.¹

Christensen's Restaurant: Two philosophers go out to dinner. Again, they fulfill the conditions for intellectual parity. When the check comes, one calculates each share as \$43, whereas the other gets to a different result, say, \$45. They are both confident that they have got it right.²

Kelly's Courtroom: You and I are members of a jury determining whether the accused is guilty. After both the prosecution and the defense rest, we have to come to a decision. We have equally good judgment and the same evidence. My verdict is guilty, yours, innocent. We are equally confident in our rightness.³

Intuitions about what one should do in cases like these vary. Some have defended the view that in such cases one should lower one's confidence in one's belief. Thus David Christensen says that in the restaurant case I should give equal credence to my belief that my share of the bill is \$43 and to my colleague's belief that it is \$45. But why should we think that? The answer, according to Christensen, is that since my peer and I are in a symmetrical position from an

¹ Richard Feldman, "Reasonable Religious Disagreements," in *Philosophers without Gods: Meditations on Atheism and the Secular*, ed. Louise Anthony (Oxford: Oxford University Press, 2007), 194-214.

² David Christensen, "Disagreement, Question-Begging, and Epistemic Self-Criticism," *Philosophers' Imprint* 11, 6 (2011): 1-22.

³ Thomas Kelly, "Disagreement and the Burdens of Judgment," in *The Epistemology of Disagreement: New Essays*, eds. David Christensen and Jennifer Lackey (Oxford: Oxford University Press, 2013), 31-53.

epistemic point of view, there is no reason for me to believe that my peer is wrong in this case apart from the disagreement itself, but our disagreement gives me no reason to demote my colleague from being my peer. For insofar as I can use our disagreement as evidence that my peer is wrong, thereby holding fast to my previous belief that my share is \$43, my peer could easily do the same. I could reasonably demote my colleague from being a peer only if I have reasons that are *independent* of my own reasoning in support of my original belief. If such reasons are lacking, there is no reason to believe that either of us has the epistemic upper hand. Rather, we should take peer disagreement as an opportunity to revise our beliefs – an opportunity, as Christensen says, *for epistemic improvement*.⁴ Conciliation is called for.⁵

A similar, but in some sense stronger position, was at one time defended by Richard Feldman. Take any of the three scenarios above and let the parties come to full disclosure of the evidence, that is, let both know that they have the same evidence. There are three doxastic options available to them: belief, disbelief, and suspension of belief. Let a *reasonable disagreement* be such that both parties to it are justified in the beliefs they hold. Feldman argues that a single body of evidence cannot justify more than one attitude towards a proposition, and since the parties have the same evidence after full disclosure, they cannot be justified in either believing or disbelieving the relevant proposition.⁶ If after full disclosure of the evidence they find themselves in disagreement, they should suspend judgment regarding the question. Since there is no *reasonable disagreement* possible after full disclosure of the evidence, the parties cannot justifiably draw different conclusions from the same evidence. Suspension of belief is the only reasonable option in cases of such disagreement.⁷

⁴ David Christensen, “Epistemology of Disagreement: The Good News,” *Philosophical Review* 116, 2 (2007): 187–217.

⁵ See also David Christensen, “Disagreement as Evidence: The Epistemology of Controversy,” *Philosophy Compass* 4, 5 (2009): 756–767.

⁶ “...a body of evidence justifies at most one proposition out of a competing set of propositions...” (Feldman, “Reasonable Religious Disagreements,” 205), and “...it cannot be that epistemic peers who have shared their evidence can reasonably come to different conclusions” (Feldman, “Reasonable Religious Disagreements,” 213).

⁷ Note that Feldman does not question the possibility of peer disagreement, only its reasonableness. We will argue that if two people are really epistemic peers, there cannot be unreasonable disagreement between them, either.

More recently, Feldman has come to defend a weaker position: what has been called the *Total Evidence View*, which is also championed by Thomas Kelly.⁸ According to this, one's response to a peer's disagreeing should be governed by one's total evidence, which consists of one's *first-order* evidence, i.e., the original evidence on which one formed one's belief, and the *higher-order* evidence consisting of the fact that one's peer disagrees. The impact and the complexity of the total evidence varies in each case, but first-order evidence still counts after full disclosure of the evidence. Thus even though there may be cases where suspension of judgment is called for, there may also be cases in which one should hold fast to one's belief.⁹ As a matter of fact, Kelly says that "there is certainly no guarantee that the uniquely reasonable response on your part is to retreat to a state of agnosticism between your original opinion and my original opinion,"¹⁰ as suggested, for instance, by Christensen's conciliationism. The reasonable response to disagreement depends on the combination of first-order evidence and the higher-order evidence "afforded by the fact that one's peers believe as they do."¹¹ What to do when a peer disagrees with you should be dictated by your judgment about your epistemic situation, and not by a general epistemic norm requiring agnosticism.

Differing as they do in their answer to how to respond to disagreement by a peer, all these writers take such disagreement to have evidential import. More than that, in the case where you believe that *p*, they take the fact that your peer disagrees with you as evidence against *p*. Thus, as mentioned above, Kelly takes the fact that one's peer disagrees with one as higher-order evidence that should be added to one's "stock of evidence;"¹² Christensen says that "the peer's disagreement gives one evidence that one has made a mistake in interpreting the original evidence, and that such evidence should diminish one's confidence in *P*;"¹³ and Feldman says that "the proposition that *S*'s peer - whose evidence concerning *p* is the same as *S*'s - disbelieves *p* is evidence against *p*."¹⁴ Therefore, in each of the

⁸ See Feldman, "Evidentialism, Higher-Order Evidence, and Disagreement," *Episteme* 6, 3 (2009): 294-312, and Kelly, "Disagreement and the Burdens of Judgment."

⁹ Cf. Feldman, "Evidentialism, Higher-Order Evidence, and Disagreement," 311.

¹⁰ Kelly, "Disagreement and the Burdens of Judgment," 35.

¹¹ Thomas Kelly, "Peer Disagreement and Higher-Order Evidence," in *Disagreement*, eds. Richard Feldman and Ted Warfield (Oxford: Oxford University Press, 2010), 142.

¹² Kelly, "Disagreement and the Burdens of Judgment," 45.

¹³ Christensen, "Disagreement as Evidence: The Epistemology of Controversy," 757.

¹⁴ Feldman, "Evidentialism, Higher-Order Evidence, and Disagreement," 298.

aforementioned cases where one believes that p , peer disagreement is taken as evidence that p is false: one should consider whether one’s belief that p is still supported by the total evidence, which now includes the evidence that one’s peer disagrees with one, and withhold belief or at least believe that p less confidently than before. And, of course, they all take it for granted that there *can* be such a thing as disagreement between epistemic peers. In section I we distinguish between strict peers and near-peers and argue that, on the account of peer-hood accepted by these and most other writers on the subject, there can be no such thing as peer disagreement. In section II we argue that in cases of disagreement between near-peers, which is, indeed, common, the near-peer’s disagreement is not evidence against the proposition one believes. It may be a reason for re-examining one’s evidence and one’s assessment of it, but that should not be taken to be the same thing as having evidence against what one believes.

I. Is Peer Disagreement Possible?

There are several slightly different characterizations of peer-hood in the recent literature. Generally, peers are characterized as having the same, or, at least, approximately the same, evidence, intellectual virtues and abilities. Thus, Feldman defines epistemic peers as being “roughly equal with respect to intelligence, reasoning powers, background information, and so on.”¹⁵ Christensen claims they are approximately equals “in terms of exposure to the evidence, intelligence, freedom from bias, etc.”¹⁶ Stewart Cohen says with respect to peers that “there is enough shared evidence that there is no reason to suppose that either party (...) is in an evidentially superior position.”¹⁷ Moreover, peers must be “in general equal in their reasoning abilities, or at least, close enough so there is no basis for supposing either party is in general the superior reasoner.”¹⁸ This does not guarantee, however, that they will *employ* their equal competence in examining the evidence equally well. That is why Christensen says that they should also “react to that evidence in the right way,”¹⁹ and Ernest Sosa emphasizes that peer-hood also depends on “how likely you are to *employ* [your general

¹⁵ Feldman, “Reasonable Religious Disagreements,” 201.

¹⁶ Christensen, “Disagreement as Evidence: The Epistemology of Controversy,” 756.

¹⁷ Stewart Cohen, “A Defense of the (Almost) Equal Rights View,” in *The Epistemology of Disagreement*, eds. Christensen and Lackey, 98-99.

¹⁸ Cohen, “A Defense of the (Almost) Equal Rights View,” 99.

¹⁹ Christensen, “Epistemology of Disagreement: The Good News,” 188.

competences].”²⁰ Gary Gutting describes peers as being equals in “intelligence, perspicacity, honesty, thoroughness, and other relevant epistemic virtues.”²¹ Kelly characterizes peers concerning a particular question as “equals with respect to their familiarity with the evidence and arguments which bear on that question,”²² and also says that “they are equals with respect to general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias.”²³ It is sometimes stressed that peers should be equals in general epistemic competence, sometimes that they must be equals with respect to the particular question they disagree about. Hence, there is a distinction between being a peer in general and being a peer concerning some proposition *p*.²⁴

However, sometimes it is stipulated that two peers must have literally the same evidence regarding *p*. According to Feldman and Warfield, we have “.... peers literally share all evidence and are equal with respect to their abilities and disposition relevant to interpreting that evidence.”²⁵ As will become clear, we think that the difference between the two characterizations – ‘roughly equal’ v. ‘literally equal’ – makes all the difference, and even though it is sometimes observed in the literature, it is usually not carefully observed. It is important to notice that both Gutting’s and Kelly’s characterizations of peers mention that they must be ‘equals,’ and not just *approximately* equals. Sometimes, this distinction is not even observed by the same author, for if we continue the passage we cited from Cohen, for instance, we have the following: “When parties to a disagreement have the same evidence and are equal in their reasoning abilities, they are epistemic peers.”²⁶

²⁰ Ernest Sosa, “The Epistemology of Disagreement,” in *Social Epistemology*, eds. Adrian Haddock, Alan Millar, and Duncan Pritchard (Oxford: Oxford University Press, 2010), 283.

²¹ Gary Gutting, *Religious Belief and Religious Skepticism* (Notre Dame: University of Notre Dame Press, 1982), 83.

²² Thomas Kelly, “The Epistemic Significance of Disagreement,” in *Oxford Studies in Epistemology*, Vol 1, eds. Tamar Szabo Gendler and John Hawthorne (Oxford: Oxford University Press, 2005), 175.

²³ Kelly, “The Epistemic Significance of Disagreement,” 175.

²⁴ Nathan King, “Disagreement: What’s the Problem? or A Good Peer is Hard to Find,” *Philosophy and Phenomenological Research* 85, 2 (2012): 249-72, is a good source for the conditions for peer-hood assumed in the recent literature.

²⁵ Richard Feldman and Ted Warfield, “Introduction,” in *Disagreement*, eds. Feldman and Warfield, 2.

²⁶ Cohen, “A Defense of the (Almost) Equal Rights View,” 99.

Of course, this may just be a loose way of putting things, but we will argue that the difference is by no means insignificant. We think we should draw a distinction between epistemic peers in the strict sense (EP) and epistemic near-peers (ENP):

- (EP) *A* and *B* are epistemic peers with respect to *p* iff they have identical evidence concerning *p*, equal cognitive abilities in general, and are equal in their employment of those cognitive abilities with respect to *p* and evidence they have.
- (ENP) *A* and *B* are epistemic near-peers with respect to *p* iff they have approximately the same evidence concerning *p*, approximately the same cognitive abilities in general, and are approximately equal in their employment of those cognitive abilities with respect to *p* and the evidence they have.

According to (EP), epistemic peers are epistemically identical with respect to *p*. Let *A* and *B* be peers, i.e., let them satisfy (EP) for some proposition *p*, and let *E_A* be *A*'s body of evidence and *E_B* *B*'s. It follows from (EP) that *E_A*=*E_B*. Now, suppose that *A* believes that *p* based on *E_A*, but *B* demurs. Then either *B* *does not* accept *E_A* or she does but does not believe that *if E_A, then p*. If the former, *E_A*≠*E_B*, contrary to the assumption. If the latter, either *A* and *B* do not have the same cognitive abilities or they are not employing those abilities in the same way. One way or the other, the mere fact that *A* and *B* disagree shows that they are not epistemic peers as defined in (EP).

Broadly speaking, there are two general conceptions of evidence. Either evidence is exhausted by one's mental states – the *internal* conception of evidence – or it includes facts about the world – the *external* conception.²⁷ We take for granted that most writers assume an internalist position, but all the above holds on either conception. Whether *E_A* and *E_B* are thought of as consisting exclusively of internal states of *A* and *B* or not, what matters for present purposes is that they be thought of as identical.

It may be objected, however, that this is too strong. Even if we ignore all the many non-epistemic differences between *A* and *B* – as we may be entitled to do here – one crucial epistemic difference may remain. Suppose *A* is a brain in a vat

²⁷ Conee and Feldman's account of evidence as occurrent thoughts is an instance of the former. See Earl Conee and Richard Feldman, *Evidentialism: Essays in Epistemology* (Oxford: Oxford University Press, 2004). The account in Timothy Williamson, *Knowledge and its Limits* (Oxford: Oxford University Press, 2000), 184-208, is one of the latter.

and *B* is not. Then, even if (EP) is satisfied, *A* and *B* will differ in that *A*'s beliefs will be false and *B*'s true. But this objection has no teeth in the present context. Saying that *A* and *B* are epistemic twins (as we will, in what follows), rather than epistemically identical makes no difference with respect to the present point. As long as the *only* epistemic difference between *A* and *B* is that their beliefs differ in truth value, disagreement between them is ruled out by (EP).²⁸

Another objection may be that non-epistemic differences cannot be ignored, since they may lead two people to evaluate *p* differently, even if they satisfy (EP). They may disagree for practical or psychological reasons. But if they disagree for non-epistemic reasons, we obviously do not have a case of *epistemic* disagreement. I may have good reasons to disagree with you about *p* when *p* is a live, forced, or momentous question for me but not for you: you are willing to accept *p* given *E*, whereas I, having more at stake, require more.²⁹ But then, again, you and I are not employing our cognitive capacities in the same way.³⁰ Such disagreements do not pose a question for the epistemology of disagreement.³¹

An arguably more epistemically relevant case is one in which the parties have different insights:

²⁸ Sosa discusses cases where disagreement is sufficient to demote the other party from being a peer. Take, for instance, a disagreement about whether you are in a certain mental state. (Sosa uses a headache as example). In this case, the only evidence I have is that you say you are. If I disagree "you would be reasonable to downgrade your opponent based essentially on your disagreement, even with no independent reason for doing so" (Sosa, "The Epistemology of Disagreement," 287). Nevertheless, Sosa points out that it is not easy to demote somebody from being a peer in "issues subject to troubling, persistent disagreement are not properly decidable in the absence of ulterior reasons" (Sosa, "The Epistemology of Disagreement," 287).

²⁹ We borrow here from William James, *The Will to Believe and Other Essays* (New York: Longmans, Green, and Co., 1897), 2-3.

³⁰ So-called attributor contextualism makes much of the difference between 'high-stakes' and 'low-stakes' situations with respect to knowledge attributions. Without endorsing that – we think, dubious – doctrine, we are deploying a similar distinction with respect to belief. (See, for example, Stewart Cohen, "Knowledge and Context," *The Journal of Philosophy* 83, 10 (1986): 574-583, and Keith DeRose, "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research* 52, 4 (1992): 913-929.) (More on the distinction between having evidence and having good reasons in section II.)

³¹ Feldman, "Reasonable Religious Disagreements," 205-6, discusses a similar case regarding disagreement resulting from opposing worldviews or starting points (see also Harvey Siegel, "Argumentation and the Epistemology of Disagreement," *Cogency* 5, 1 (2013): 135-170, for a discussion of "deep" disagreement). For a much earlier discussion on this, see Robert Fogelin, "The Logic of Deep Disagreements," *Informal Logic* 7, 1 (1985): 3-11.

Each may have his or her own special insight or sense of obviousness. But each knows about the other’s insight. Each knows that this insight has evidential force. And now I see no basis for either of them justifying his own belief simply because the one insight happens to occur inside of him. A point about evidence that plays a role here is this: evidence of evidence is evidence. More carefully, evidence that there is evidence for P is evidence for P. Knowing that the other has an insight provides each of them with evidence.³²

But here, again, if *A* and *B* have different insights with different evidential force, they do not satisfy (EP). If you have a revelation or vision that I do not have, we have different evidence and are thus not peers. If you are convinced by the argument from design and I am not, we differ in how we assess the evidence we both have and are thus not peers.³³

Catherine Elgin has argued that a definition of peer-hood such as (EP) is too narrow and that for that reason it is not helpful in understanding *real-world* disagreement. She says that

in the recent debates about disagreement ‘epistemic peer’ is defined quite narrowly. It requires having *the same* evidence, and reasoning abilities. So it is not surprising if an ordinary person lacks epistemic peers with respect to a particular, mundane issue. If Jen and Jon have even slightly different relevant reasoning abilities or evidence pertaining to the causes of the Civil War, they are not epistemic peers with respect to the subject. Given the vicissitudes of education and abilities, and the idiosyncrasies of evidence gathering, ordinary epistemic agents are apt to have few epistemic peers.^{34, 35}

Elgin is correct in observing that peer-hood as defined in (EP) is an idealization, and that we need a more generous account of it to deal with real-world cases of disagreement, one on which two disagreeing parties have only “pretty much the same evidence, reasoning powers, training, and background information.”³⁶ This echoes Feldman’s requirement that they be “roughly equal,”

³² Feldman, “Reasonable Religious Disagreements,” 208.

³³ The same goes for Conee’s suggestion about different intuitions. See Earl Conee, “Rational Disagreement Defended,” in *Disagreement*, eds. Feldman and Warfield, 69-90.

³⁴ Catherine Elgin, “Persistent Disagreement,” in Feldman and Warfield, *Disagreement*, 56-57.

³⁵ Similar doubts about the possibility of strict peer-hood are voiced by Jonathan Matheson, “Disagreement: Idealized and Everyday,” in *The Ethics of Belief: Individual and Social*, eds. Jonathan Matheson and Rico Vitz (New York: Oxford University Press, 2014), 315-330, King, “Disagreement: What’s the Problem?” and Siegel, “Argumentation and the Epistemology of Disagreement.”

³⁶ Elgin, “Persistent Disagreement,” 57.

which we believe is captured by (ENP). In section II we will discuss cases involving epistemic near-peers, which may capture the more ‘generous account’ that Elgin has in mind. However, Elgin’s criticism of the stricter notion of peer-hood is very different from ours. We do not say that (EP) should be abandoned in favour of a more generous account, given its inapplicability in the real world. Our claim is, rather, that peer disagreement in the *ideal case* is impossible. Elgin does not seem to realize this, for she goes on to say that her solution to cases of near-peer disagreement also works if the “standard conception of an epistemic peer is used.”³⁷ And, by “standard” here she means (EP).³⁸

Thus we agree that genuine peer disagreement in the strict sense of peer-hood is problematic, but not because it is rare or hard to come by. We maintain that if peer disagreement is understood as in (EP), there *could not be* such a thing. There can, of course, be disagreement between *apparent* epistemic peers, who appear to be such because they satisfy (ENP) and are, in Feldman’s words, *roughly* equal. This is also what Elgin’s “generous” characterization of peers comes to. Whatever one calls it, it amounts to a more permissive definition of peer-hood, taking into account at least some differences in the evidence or the intellectual capacities (or both) of the disagreeing parties. Thus, you may *appear* to be my peer as defined in (EP) and yet disagree with me. Indeed, such cases are common. To appear to be my peer, you must be my near-peer; if you were not, it would very likely be obvious that you are not. In such a case, if I understand peer-hood *a la* (EP), I know that you cannot be my peer. What I do not know is whether you are my epistemic superior or my epistemic inferior, that is to say, whether you are more or less competent or have stronger or weaker evidence or both.³⁹ If I believed you to be my superior, I would have reason to defer, if to be my inferior, to stand firm. So far I have no reason for doing either; thus Feldman is right that suspension of judgment is called for. However, to agree that suspension of judgment is called for *pro tem* is not to agree that that is all that is called for. Other, obvious, things to do are to re-examine both bodies of evidence and to look for additional evidence.

³⁷ Elgin, “Persistent Disagreement,” 57.

³⁸ It is not clear whether one should regard (EP) or (ENP) the “standard” conception of peer-hood – if, indeed, there is one at all. We take this to be at least partly due to the back and forth between ‘equal’ and ‘approximately equal’ in the definitions of peer-hood, and also to a certain looseness or lack of precision in many definitions, as illustrated above. In any case, we use ‘peer’ for the strict case, (EP), and ‘near-peer’ for the notion defined in (ENP).

³⁹ This is a broad definition of epistemic superiority. Some may prefer a narrower one, focusing only on better or superior cognitive ability. This makes no difference to the present argument.

We know that we cannot both be right (though we could both be wrong). We therefore know that one of three things must be true. One of us may be misjudging the evidence he has, the evidence one of us has may be misleading, or there is other evidence to be found that would tip the scales one way or the other.

I may not know which of us is epistemically superior, but your disagreeing gives me reason to try to find out, and what better way than to ask you. Suppose you claim to have evidence I do not have. If I believe you and think you my epistemic equal in other respects, I should defer, if I do not believe you or think you epistemically inferior, I should not. Whichever the case, I do not take myself to be deferring to, or refusing to defer to, an epistemic peer. Suppose you claim that the evidence we both have shows something different than what I think it does. I ask you to explain and either accept your explanation or do not. Doing the first is tantamount to deeming you my epistemic superior, doing the second, my inferior. Again, I do not take myself to be deferring to, or refusing to defer to, an epistemic peer. Whatever I do, I do precisely because I do not take us to be the epistemic peers we may have appeared to be – the mere fact that we disagree is proof that I am right in this.

Thus in Christensen’s example, we both know that (at least) one of us is mistaken. Given that our evidence is obviously the same, it follows that we are not equal in our cognitive abilities – or, at least, in our employment of them with respect to the matter in question – and that we both know this. If each one of us is confident that he is right, as Christensen says, that is tantamount to his thinking that the other is his epistemic inferior. However, such confidence is compatible with thinking that it may, after all, turn out to be the other way around. Hence the only sensible thing to do is to set about finding out which of us is mistaken. This is, obviously, what two rational people finding themselves in such a situation would do: they would go over the bill together. Similar remarks may be made about Feldman’s case (“Fetch the binoculars!”). Disagreement of this sort is, indeed, occasion for epistemic improvement. But that is not the same thing as “giving equal credence” to a proposition one does not believe and for which one’s only evidence is that someone who may or may not be one’s epistemic superior believes it.

Distinguishing between higher-order evidence, as do Kelly and (now) Feldman, makes no difference to any of this. If my total evidence includes evidence that you have come to a different belief on the basis of first-order evidence we both have, we do not satisfy (EP) with respect to parity of cognitive ability. Furthermore, (EP) cannot be satisfied with respect to total evidence so understood.

You and I cannot have the same total evidence, as mine will include evidence of your disagreeing with me and yours will include evidence of my disagreeing with you.⁴⁰

II. Near-Peer Disagreement and Evidence of Evidence

a) Evidence of What?

While there can be no such thing as peer disagreement strictly speaking, cases of near-peer disagreement are not uncommon. We have suggested that in such cases there is no need to look for any special principle telling the near-peers how they should react to their disagreement: general principles of sensible epistemic conduct are all they need to employ. Some, however, take the fact that a near-peer disagrees with one as itself evidence against one's position and argue that it is sufficient for giving up one's position, either suspending belief (Feldman) or going as far as deferring to one's disputant (Christensen).

Feldman takes the following principle to be the "key evidential fact about disagreements:"

(KEF) The proposition that *S*s peer - whose evidence concerning *p* is the same as *S*s - disbelieves *p* is evidence against *p*.^{41, 42}

The first problem with (KEF) is that it does not capture the idea behind the slogan 'evidence of evidence is evidence,' as it is presumably intended to do.⁴³ If my

⁴⁰ Our total evidence can be the same if we ignore the indexical aspect and focus on the character of the sentence describing the higher-order evidence, rather than on its content. If the evidence is described only as one's peer's disagreeing, we may both be said to have it. But then if we have the same total evidence and still disagree, that must be because we are not equals in cognitive ability.

⁴¹ Feldman, "Evidentialism, Higher-Order Evidence, and Disagreement," 298.

⁴² Disbelieving a proposition may be taken to mean not believing it or believing its negation. Thus one's peer may disagree with one either by denying what one believes or by merely refusing to affirm it. Feldman's 'evidence against *p*' suggests that he has the former in mind. It is not clear whether he intends the principle to apply in the latter case. (In what follows we will mean by 'peer' 'near-peer' as defined in (ENP).

⁴³ Feldman has different formulations of the principle 'evidence of evidence is evidence' in "Reasonable Religious Disagreements," 208, "Evidentialism, Higher-Order Evidence, and Disagreement," 308, and in Richard Feldman, "Evidence of Evidence is Evidence," in *The Ethics of Belief*, eds. Matheson and Vitz, 292. (KEF) is supposed to be a principle about epistemic support (cf. Feldman, "Evidentialism, Higher-Order Evidence, and Disagreement," 298),

peer’s evidence is really the same as mine, our disagreement can be put down only to superior cognitive ability on the part of one or other of us. Rational disagreement can be explained only by one party’s having either more evidence or better judgment. (KEF) explicitly rules out the former, so only the latter can be in play. Then the only thing of which my peer’s disagreeing with me can be evidence of is that one of us has better judgment than the other. But that is not evidence that *his* judgment is better than mine.

In any case, what is presumably intended is that in having evidence that you do not believe what I believe I have evidence that you have evidence against what I believe – rather than that your judgment is better than mine – and *that* gives me evidence that what I believe is false. As already noted, as long as we differ in either our evidence or in our cognitive ability, we cannot be peers strictly speaking. Since disagreement between peers is impossible, our disagreement shows that we are not. (KEF), if taken to be about strict peers, is incoherent, requiring as it does thinking of the parties as both peers and not peers.

A recognition of this may lie behind Feldman’s characterization of epistemic peers as “*roughly* equal” in their evidence and cognitive abilities; in effect, as only near-peers.⁴⁴ However, insisting, as we have repeatedly, that any difference in what evidence two people possess or in their cognitive ability entails that they are not peers is no mere pedantry. We will argue that if one’s disagreement is with someone who is only roughly one’s peer, no special problem arises. More investigation, concerning the evidence the other has or his cognitive abilities (or both) is all that is required.

Feldman thinks otherwise. He thinks that in cases of what we are calling near-peer disagreement the following principle of evidence holds:

(EEE) Evidence that *Ss* near-peer disbelieves *p* is evidence against *p*.

specifically in cases of peer disagreement. We believe there is an ambiguity in Feldman’s different formulations. While Branden Fitelson, “Evidence of Evidence is not (Necessarily) Evidence,” *Analysis* 72, 1 (2012): 85–88, and Juan Comesaña and Eyal Tal, “Evidence of Evidence is Evidence (Trivially),” *Analysis* 75, 4 (2015): 557–559, discuss (EEE) assuming that the evidence one has is evidence of *what* the evidence the other has is (or is not, in Fitelson’s case), we take it to mean only that it is evidence that the other has *some* (or something he takes to be – see below) evidence. For more on this, see Fabio Lampert and John Biro, “What is Evidence of Evidence Evidence of?” *Logos & Episteme* VIII, 2 (2017): 195–206.

⁴⁴ Feldman, “Reasonable Religious Disagreements,” 201, our emphasis.

Thus when one learns that one's near-peer does not believe what one does, one acquires evidence against what one believes.

One problem with this is that the evidence I have for the proposition that you do not believe that p is evidence only that you take yourself to have evidence that p is false.⁴⁵ Even supposing that my evidence that you do is good evidence that you do (which it may not be), the truth of *that* proposition is compatible not only with p 's being true but with my not having any evidence that p is false. Suppose I believe that the accused is innocent and you say that you believe him guilty. Assuming that you are sincere and you believe what you do because you take yourself to have evidence of his guilt, is the evidence I have – that you take yourself to have such evidence – evidence that the accused is guilty? Surely, only if I think that the evidence you take yourself to have is *good* evidence. If the evidence in question is misleading, your having it gives me no reason to believe what you do.

It is one thing to have evidence that p and another to have evidence that is good evidence that p . Since 'evidence that p ' may be taken to mean either, for all I know, what you have might be misleading evidence. If that is the case, the fact that you disbelieve p does not speak against p . To decide whether it does so, I must find out what your evidence is and weigh it myself to the best of my ability. If I decide that your evidence is good evidence, that is tantamount to deciding that, in coming to believe what you did, you were my epistemic superior. Now that I have the evidence you have, we are again epistemic peers – but given that I think the evidence in question is good evidence, we no longer disagree.

We should thus distinguish between two senses of 'evidence,' the *neutral* sense and the *probative* sense, as we may dub them. Accordingly, ' S has evidence that p ' may mean one of two things. It may mean that E (where E is some proposition, state of affairs, fact – it does not matter which) is taken by S to speak in favour of p . In this sense, something is evidence if it is deemed to be so, and to say that S has evidence that p is not to say that what he has is evidence that supports p . In the latter, probative sense, to say that S has evidence that p is to say that S has evidence that *does* speak in favour of p .

In which of these two senses should we understand 'evidence' as it occurs in (EEE)? If in the first, the principle asserts merely that A has evidence (in one or the

⁴⁵ Or that you see the evidence we both have as evidence for $\neg p$, rather than for p . For now, we ignore this possibility, though we will come back to it later. In any case, part of the assumption is that S 's near-peer disbelieves p because of what she deems to be evidence against it.

other sense) that *B* takes himself to have evidence that speaks in favour of *p*. But that should be of interest to *A* only if he has reason to believe that *B* is his epistemic superior.

Suppose that *A*’s evidence is *B*’s saying that he has evidence and *A* takes that as evidence that *B* has probative evidence that *p*. He would be mistaken if *B* were lying. Even if *B* is truthful, *he* may be mistaken: the evidence he takes to be probative may not be. In neither case does *A* have probative evidence that *p*. For (EEE) to have the bite it is thought to have, the second occurrence of ‘evidence’ must be understood in the probative sense. The trouble is that *A* (and we) have no grounds for thinking that the evidence *B* has is probative unless we know what that evidence is and judge it to be such. His disagreeing does not by itself constitute such grounds.

Let us label the two occurrences of ‘evidence’ in (EEE) as ‘evidence₁’ and ‘evidence₂.’⁴⁶ What we are interested in is whether someone asserting (EEE) is right in claiming, as its proponents appear to, that ‘evidence₂’ is probative, that it speaks in favour of, or against, the proposition in dispute.

Imagine that we learn from a reliable source that a respected colleague has been claiming that LBJ was behind the assassination of JFK. We have evidence that he takes himself to have good evidence for the claim. Until we see what that evidence is, however, we have good reason to think him a whacko. Our evidence that he takes himself to have good evidence clearly does not mean that we have good evidence, not even that we have what we (perhaps mistakenly) take to be good evidence. Suppose that we learn *what* he takes to be evidence, and it *is* whacky; now we have good evidence that he *is* a whacko. Suppose, on the other hand, that we find the evidence he takes to be good evidence. Now we indeed have (what we take to be good) evidence ourselves. But this does not help (EEE). The evidence we now have that LBJ was involved is more than just that our colleague has what he takes to be good evidence, which is all we had before and which (EEE) says is sufficient for us to have evidence – indeed, good evidence – that LBJ was involved.⁴⁷ Thus it is a mistake to say that “... mere disagreement... is in general a

⁴⁶ Following (loosely) Richard Feldman, “Evidence of Evidence is Evidence,” *Keynote lecture at Feldmania: A Conference in Honor of Richard Feldman*, UT San Antonio, 2011, and Fitelson, “Evidence of Evidence is not (Necessarily) Evidence.” As a slogan, (EEE) has three occurrences of ‘evidence:’ evidence₁ of evidence₂ is evidence₃. But only two appear in the principle as formulated above.

⁴⁷ This is why eyewitness testimony is admissible as evidence while hearsay is not. The difference

perfectly legitimate piece of evidence ... and so often a good enough reason to demote [someone] from the status of a peer.”⁴⁸

b) Evidence and Content

Certainly, your not believing that *p* gives me reason to re-examine the evidence on which I base my belief that *p* and my assessment of that evidence. And, if I have reason to think that you may be my epistemic superior, it may even give me reason to defer to you, pending – or lacking the opportunity for – such re-examination. Relying on experts is generally good policy and often unavoidable. But believing something on mere authority is not the same thing as believing something on the basis of evidence. For the latter, there must be what we may call a *content-connection* between one’s evidence/ground and one’s belief. Roughly, one must think that *p* is true *because* one thinks that what is adduced as evidence for it is true and is evidence that speaks in favour of *p*. That connection is missing when one bases a belief on mere authority, however reasonable and prudent doing so may be. If I think that *p* is true because I trust the expert who declares it to be so, I do not, Euthyphro-like, think that it is his so declaring it that makes it true. But that he so declares is all Feldman’s slogan gives us, even when ‘evidence?’ is read as ‘good evidence’ (indeed, even if it is read factively, as entailing *p*). Put simply, even if my believing that you have good evidence that *p* is false is enough to give me reason to believe that *p* is false, neither my so believing nor your actually having what I believe you to have is evidence for me that *p* is false. To have that, I have to

between the two is not that the former is necessarily more reliable than the latter. We may trust the report of a usually reliable source more than we do the eyewitness testimony of one we suspect of inattention or perjury. But the eyewitness testimony is given under oath, whereas what the reliable source says is not (even if the report that he said it is). Thus, speaking *de jure*, the veracity of the former is guaranteed in a way that of the latter is not. Hence only the former is deemed probative evidence. (Of course, *de facto*, perjury is always possible. But that does not affect the point.)

⁴⁸ David Enoch, “Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement,” *Mind* 119, 476 (2010): 981. Enoch argues that there is no general recipe for how to respond to peer disagreement and that, partly for that reason, the notion is of less interest than is generally thought. While we have argued that, strictly speaking, there is no such thing as peer disagreement, we agree with both these points with respect to near-peer disagreement, even though our reasons are very different from his.

know what your evidence is and see it as probative of p .⁴⁹ Unless I have done so, you are still the sole possessor of the evidence in question.

Considerations concerning explanatoriness may help here. If I have good evidence that p , the evidence I have is explained (albeit defeasibly) by p 's being true. My having it is not so explained.⁵⁰ My having evidence that you disbelieve p is not explained by p 's being false, even if you do have the evidence you claim to have and even if that evidence is good evidence that p is false. Even if you tell me what your evidence is, and even if that evidence is explained by p 's being false and I see that it is, the *explanans* of my having the evidence I have is your telling me that you do, not p 's being false. The only evidence (EEE) allows me is evidence of the fact of your disagreement. *That* is not explained by p 's being false, even if the evidence you have is. And I am in no position to know whether it is, just because you say that it is. I may or may not believe you. If I do, I may well defer to you, but that would be for the substantive reason that, thinking that you are my epistemic superior, I think your evidence likely to be good, not because the concept of evidence demands it. The lack of an explanatory connection is a reflection of a lack

⁴⁹ Some may take this distinction to be between *direct* and *indirect* or *prima facie* evidence, rather than evidence and good reason. (See, for instance, Siegel, "Argumentation and the Epistemology of Disagreement," 144-145, fn. 1.) We believe that there is no such thing as indirect evidence. However, suppose there were. Our case against (EEE) would still hold. Suppose we modify (EEE) as

(EEE)* Evidence that S 's near-peer disbelieves that p is *indirect* evidence against p .

The problem with (EEE)* is that it is too weak. It does not matter whether you call it *indirect* or *prima facie* evidence, or even just (good) reason, as we did. If all that near-peer disagreement gives us is *indirect* evidence, then there is nothing substantive about (EEE)*. As a matter of fact, (EEE)* is a truism! Nobody has ever denied that a near-peer's disagreement gives one reason (or, *indirect* evidence) to pause and re-examine one's evidence. To deny that would be tantamount to endorsing a full-blooded dogmatism. If that is all the *Total Evidence View* has to say, we have no objection. But to defend the claim that we are no longer justified in our belief after our near-peer disagrees with us we need something much stronger, that is, we need evidence that, once added to our original body of evidence, makes it inconsistent. We need (EEE), rather than (EEE)*.

⁵⁰ The pavement's being wet is explained by its having rained; my seeing that the pavement is wet is explained not by that but by my looking at it, etc. It is not that its having rained may not be part of the total explanation, as may the work of the night street-cleaning crew. But whichever explains the pavement's being wet, it is not what explains my seeing that it is.

of what we have called a content-connection between the evidence I have about your attitude to p and p .⁵¹

Let us suppose that my evidence that you do not believe what I believe is that you tell me so. First, I have no way to tell whether you are telling the truth. More important, if you say that you do not believe as I do because you have evidence against what I believe, I have no way to tell whether the evidence you claim to have really *is* evidence against what I believe. If I found out what it was, I may not see it as evidence against what I believe. Of course, it is possible that you would be right and I wrong. But for the evidence that you have evidence against my belief to be something I see as evidence against my belief (and thus as relevant to my epistemic conduct, as (EEE) is supposed to be) you must really have evidence that I see as evidence against it. Your saying that you do is not by itself evidence that you do; hence I do not *thereby* have evidence against my belief. This is so even if it is probable that you are right and even if I think that you are right. Even if the latter gives me reason to defer, that does not mean that I have been given evidence against the proposition I believed before discovering your belief concerning it.

What is missing here is precisely what is guaranteed (*de jure*) with eye-witness testimony. Being under oath, a witness can be assumed to be truthful in a way that the source of hearsay cannot. This means that a witness' saying that he saw the accused commit the crime is as good as my seeing it and thus counts as evidence of his guilt. His reporting that someone said that he witnessed the crime is not. In other words, we assume that there is a content-connection between the eye-witness testimony that p and p .⁵²

⁵¹ Take the congressman investigated for influence-peddling. He is asked whether he was ever at the hotel where he is alleged to have received the bribe. Knowing that he is likely to have been seen leaving the hotel more than once, he takes the fifth. Evidence that he has something to hide, certainly: but it is philandering, rather than corruption. It would be hasty to conclude that the evidence we have that he has some evidence concerning the allegation is evidence that he is guilty, when the evidence he has in fact exonerates him of that charge.

⁵² Fitelson, in offering counter-examples to various ways of understanding the claim that evidence₁ of evidence₂ is evidence₃, seems to recognize the lack of content-connection we have been pointing to. He describes cases in which evidence₁ is epistemically irrelevant to p in that it does not raise p 's epistemic probability, thus "... nothing p -supporting [i]s part of" the evidence₁ one has. We agree, but we are puzzled by the parenthetical 'necessarily' in the title of Fitelson's paper. We think such considerations show that evidence₁ in and of itself is *never* evidence₂ that p in cases disagreement. For more on Fitelson's case, see Lampert and Biro, "What is Evidence of

Suppose that I have probative evidence₁ that my near-peer thinks he has probative evidence that *p* is false. According to (EEE), this entails that I have probative evidence₂ that *p* is false. This can be the case only if his evidence *is* probative. In order for me to think that it is I must either know what evidence₂ is and take it to be evidence that *p* is false or take my near-peer to be my epistemic superior. If I do the latter, that does give me reason to believe that *p* is false. But it does not give me evidence that it is. Even if I do take my near-peer to be my epistemic superior and thus regard his evidence as probative, that still does not entail that I have evidence₂ that *p* is false. Given that I do not know what his evidence is, my belief is based on mere authority. And while that, as we have allowed, may give me a good reason to believe that *p* is false, it does not give me (probative) evidence that it is for want of a content-connection between the evidence I have (evidence₁) and *p*. On the other hand, suppose that I know what his evidence is and take it to be good evidence that *p* is false. Now it can be said that I do have evidence that *p* is false, except that it is not evidence₂, as it does not flow from evidence₁, that is, from the fact that I have evidence that my near-peer has evidence that *p* is false. If I know what his evidence is and judge it to be good, that is because I see a content-connection between it and *p* and thus believe that *p* is false because his evidence is as it is. However, its being as it is is obviously not entailed by my near-peer’s thinking that it is. Thus (EEE) fails either way, whether I take my near-peer to be my superior or not.

As we have already emphasized, this is not to say that your telling me that you have evidence against what I believe cannot, even by itself, be reason for changing my mind. If I think you trustworthy, and examining the evidence you claim to have would be impractical, that may be the reasonable thing to do. But none of this is of any help to (EEE). That I have reason to suspend judgment, or even defer to you, in the face of evidence that you have evidence against a proposition I believed before gaining my evidence does not mean that I have evidence against the proposition. Consider another case of Feldman’s, that of the two detectives, one of whom has what he believes is evidence incriminating Righty, with the other thinking that he has evidence convicting Lefty. Suppose that the evidence on the basis of which the first concludes that Righty is the culprit was planted by Tricky. We can concede that that evidence, were it genuine, would (tend to) prove Righty guilty and also that our first detective, having no reason to suspect a frame-up, is justified in believing that he is. We would draw the

Evidence Evidence of?”.

same conclusion in his place. This does not make it true that the evidence he has is actually evidence that Righty is guilty and if we know that it is a plant we do not so regard it.⁵³ Should the second detective do so on finding that the first has what he takes to be evidence against Righty? Since he does not know – any more than does the first detective – that the evidence is phony and has no reason to suspect that it may be, the right thing for him may well be to suspend judgment, at least pending further investigation. But that does not mean that he now has evidence that Righty is guilty. Furthermore, even if Righty *were* guilty, his being so, even if it explained the first detective's evidence, would not explain the evidence the second detective has about the first.

Two central facts about what we are calling probative evidence are, first, that no evidence can be such evidence for me if I do not know what it is and, second, that someone's (even an expert's) thinking that something is the case is not evidence that it is the case. (EEE) fails to respect both these facts. First, it asserts that your having some evidence, I know not what, counts as evidence for me. Second, it says not just that your thinking that something is the case is evidence that it is but something even stronger, namely, that your thinking that you have evidence that something is the case is evidence that it is. In fact, one's having evidence and one's thinking that one has evidence are independent of each other: not only is it possible to think that one has evidence that *p* and fail to have it, it is also possible to have evidence that *p* and fail to realize that one has it.

c) Evidence and Having Evidence

Tal and Comesaña have recently argued that while (EEE) is false on most understandings of it (including Feldman's own), there is one interpretation on which it is true, but trivial.⁵⁴ Their criticisms of other interpretations, as well as their defence of their own, rests on two distinctions they make. The first is that between someone's *having* evidence and there *being* evidence (which, perhaps, no-one has). The second is between *de re* and *de dicto* readings of (EEE).⁵⁵ Tal and

⁵³ We have evidence – Shakespeare says so – that Othello sees the handkerchief he had given Desdemona in Cassio's hands and takes that as evidence of her infidelity. Were we not aware of Iago's machinations, we might, too. As it is, we know that it is not.

⁵⁴ See Eyal Tal and Juan Comesaña, "Is Evidence of Evidence Evidence?" *Nous* 51, 1 (2017): 95–112.

⁵⁵ The relevant version of (EEE), according to Tal and Comesaña, is the following:

(Existential EEE1 *de dicto*): $\forall(e)\forall(p)\forall(S)\forall(\alpha > 0)\forall(\beta > 0)((F(e, \exists(e')(T(e')) \wedge$

Comesaña argue that (EEE) is true (only) if interpreted as saying that my having evidence that you have evidence that p is evidence that there is something that is evidence that p . This does not entail that I have evidence that p and thus does not require that I know what your evidence is.

As far as the first of these distinctions goes, we agree that there being evidence does not entail that someone has it.⁵⁶ But someone’s having evidence does entail there being evidence unless we read the first non-probatively. But even if we assume that I have probative evidence that you have evidence that p , my evidence shows that there is evidence that p in the probative sense only if the evidence you have is probative. Tal and Comesaña assume that it is.⁵⁷ But, as we have already argued, to have evidence that there is probative evidence one has to know what the evidence supposedly probative is and to judge it probative. Doing this requires knowing what the evidence in question is, rather than just have evidence that, whatever it is, it exists, which is all (EEE) allows. Knowing what the evidence is is tantamount to having it.

$$F(e', p, \alpha), \beta) \rightarrow \exists(\gamma > 0)(F(e, p, \gamma)))$$

which is to be read as “for all e and p , if e is evidence that there is evidence e' for p , then e is evidence for p ” (Tal and Comesaña, “Is Evidence of Evidence,” 102). Their own version, however, is this:

$$\begin{aligned} &(\text{Existential EEEI } de \text{ dicto no defeat: } \forall(e)\forall(p)\forall(\alpha > 0)\forall(\beta > 0)\forall(\gamma > 0) \\ & (F(e, \exists(e')(T(e') \wedge F(e', p, \alpha)), \beta) \wedge (F(e \wedge \exists(e')(T(e') \wedge F(e', p, \alpha), p, \gamma)) \rightarrow \exists(\delta > 0) \\ & (F(e, p, \delta))) \end{aligned}$$

which is to be read as “for all e and p , if (i) e is evidence that there is evidence for p and (ii) e is not a defeater for the support that the proposition that there is evidence for p provides for p , then e is evidence for p ” (Tal and Comesaña, “Is Evidence of Evidence,” 108). As will become clear, our criticism applies to both versions.

⁵⁶ And one can even have evidence that there is evidence that p without having evidence that someone has evidence that p . Suppose I have (non-probative) evidence that you, having killed your shipmate, threw the murder weapon into the sea. This is evidence that there is evidence that you are guilty; alas, no-one has it or, probably, ever will.

⁵⁷ Actually, they make it clear that they think of evidence as factive: “We will assume here that someone has a proposition as evidence only if that proposition is true... To make the factivity transparent, we will symbolize that subject S has evidence E with ‘ $T(e) \wedge S(e)$.’” (4) If ‘evidence’ in both ‘ S has evidence’ and ‘there is evidence’ is factive, the first clearly entails the second. (EEE) may be true interpreted as saying that if one has evidence that entails that one’s peer has evidence that entails that p one has evidence that p . But since one never has such evidence, so interpreted it is useless as a guide to epistemic conduct.

Tal and Comesaña say that

... evidence that there is evidence for p may be sufficient for counting as evidence for p even if nobody has it (suppose, for instance, that all the researchers involved in studying whether p declare that they found excellent evidence for p , but then die before telling us what the evidence is).⁵⁸

But “declare that they found evidence” does not mean that they have probative evidence, however excellent they *say* it is. It does not even entail that there *is* evidence probative *or* non-probative. The researchers may be mistaken in thinking that what they have is evidence even in the non-probative sense. Suppose the value of neither one of two variables they think is evidence of their hypothesis really is such., They may be mis-reading a measurement, taking it to give the value of one of the variables when in fact it gives that of the other. Then they do not even have non-probative evidence, even though they think they do.

Regarding the second distinction, between *de re* and *de dicto* readings of (EEE), Tal and Comesaña say the following:

Just as there is a difference between believing that a specific person is a spy and believing the existential proposition that there are spies, there is an analogous difference between e being evidence for a specific proposition which is evidence for p , and e being evidence for the existential proposition that there is evidence for p .⁵⁹

On the *de re* reading, (EEE) entails that the one for whom evidence₁ of evidence₂ is supposed to be evidence₃ knows what evidence₂ is. If so, it is not the evidence that someone else has it or that it exists that is evidence that p but evidence₂ itself. On this reading, in having E_1 one *eo ipso* has E_2 , making the first E of (EEE) idle. On the other hand, since on the *de dicto* reading (which Tal and Comesaña argue is the correct one) it is not specified what the evidence that p is,

⁵⁸ Tal and Comesaña, “Is Evidence of Evidence,” 103. We have avoided the locution ‘evidence for p ’ preferring ‘evidence that p ,’ as substituting for ‘ p ’ yields ‘evidence that it is raining’ in the latter but the nonsensical ‘evidence for it is raining’ in the former.

⁵⁹ Tal and Comesaña, “Is Evidence of Evidence,” 97. They offer the following formalizations for the two readings:

(e is evidence that there is *de re* evidence for p): $\exists(e')\exists(\alpha > 0)\exists(\beta > 0)(F(e, e', \alpha) \wedge F(e', p, \beta))$.

(e is evidence that there is *de dicto* evidence for p): $\exists(\alpha > 0)\exists(\beta > 0)(F(e, \exists(e'))(T(e') \wedge F(e', p, \alpha)), \beta)$.

one cannot judge whether it is probative. Not being able to do that means that one does not have evidence that *p*. It *may* be the case that there *is* evidence that *p* and that the experts had it. There may also be, as Tal and Comesaña suggest, such a thing as evidence that no one has. Neither is any help to (EEE). The experts’ declaration may, of course, give others excellent *reason* to accept *p*. But it does not give them evidence that *p*.

Not only does evidence that someone has evidence that *p* fall short of being evidence that *p*, sometimes it stands in the way of having evidence that *p*. Think about Lois Lerner’s taking the fifth at the recent congressional hearings on the IRS’ allegedly targeting conservative organizations. Her doing so is evidence that (she thinks) there is evidence to support the allegations. Yet not only does the committee not have such evidence as a result of her declining to answer questions, her exercising the right to do so is intended (by both her and the Constitution) to make sure that it does not. The committee – and we – have ample reason to be suspicious. But that is not the same thing as having evidence, which is why the committee is demanding to see Lerner’s e-mail messages and she is refusing to let them.

What about Tal and Comesaña’s distinction between having evidence and there being evidence, though? Is her refusal not evidence that there *is* evidence against her, even if it is shielded from the committee? Well, no. She may think that there is and be mistaken. She may be badly advised by her attorney. Once again, the evidence we have that she (thinks) she has incriminating evidence may be a reason for thinking that she does, but it is not evidence that she does. (As President Obama so intriguingly put it, “there is not a shred of evidence” of wrong-doing.)⁶⁰

Thus (EEE) understood along the lines Tal and Comesaña recommend is as vulnerable to these objections as on any other interpretation.

Conclusion

Disagreement between people who are peers strictly speaking is impossible. Disagreement between people who are peers to all appearance is frequent, but we know that the appearance is deceptive. Of course, knowing that one party must be epistemically superior is one thing, knowing which is another. While the latter can

⁶⁰ Nothing hangs on the particulars of this case in its involving constitutional considerations. Hilary Clinton’s refusal to make her private server available for inspection by an independent party may be seen as evidence that she has something to hide, but – happily or unhappily, take your pick – it is not evidence that she does, as she well knows.

be hard to figure out, trying to do so is usually the proper response to disagreement. When that is impractical, we can choose between deferring to expertise, suspending judgment and agreeing to disagree, depending on the case and, in particular, on the importance and urgency of the matter. In whichever of these ways we respond, general principles of sensible epistemic conduct are all we need to rely on. We do not need (EEE) in addition to these.

Returning now to the question of dependence. If in the face of a near-peer's disagreement all one needs are the general principles of sensible belief-formation, with the near-peer's belief playing no evidential role, it cannot be said that in forming the belief one does on learning of the disagreement one is epistemically dependent. It is instructive to compare the situation with the genuine case of epistemic dependence we have when in forming our belief about something we take into account the opinion of an expert. We believe that the expert has evidence we do not have. The evidence we have of that is the independent evidence we have that he is an expert where we are not. This is just what is missing in the case of near-peer disagreement. What alerts us to the fact that someone who appears to be our peer is not is the fact that he disagrees with us. That tells us that he is either our epistemic superior or our epistemic inferior, but it does not tell us which. It does not tell us even whether our epistemic difference consists in having different evidence, rather than differing assessment of our common evidence. But even if it happens that it is the former, we have no reason to believe that any evidence our disputant may have that we do not is probative evidence, something we assume with someone we believe, on good grounds, to be an expert. This is why we take his evidence into account in a way we have no reason to do with a near-peer's. As we have argued, even if we have evidence that a near-peer has evidence against our belief, that gives us no evidence against our belief. Contrary to what (EEE) implies, we are not epistemically dependent on a near-peer who disagrees with us.⁶¹

⁶¹ We are grateful to Harvey Siegel, Pedro Merluzzi, and Jonathan Matheson for many helpful comments on previous drafts of this paper.

AN AXIOM LINKING NECESSITY AND OBLIGATION PROVIDED BY PRIOR AND ITS ANALYSIS UNDER CARNAP'S METHOD

Miguel LÓPEZ-ASTORGA

ABSTRACT: Although written long before, in 2012 a work by Prior presenting a system that was able to demonstrate Hintikka's theorem was published. Maybe one of the most relevant elements of that system is an axiom that clearly relates necessity, and hence modal logic, to obligation, and hence deontic logic. This paper analyzes that axiom based upon Carnap's method of extension and intension in order to show that it should be accepted. Thus, the paper is intended to give further evidence supporting not only the aforementioned axiom, but also Prior's system in general and, accordingly, Hintikka's theorem.

KEYWORDS: Rudolf Carnap, L-concepts, necessity, obligation, Arthur N. Prior

Introduction

There is no doubt that the theorem known as 'Hintikka's theorem' is, at least, disturbing. The reason for this is that the theorem provides that, if something is not possible, it cannot be allowed, an idea that, in principle, seems to be very rare. However, Prior thought about a system that proved it. Although probably written some decades before,¹ the paper² was not published until 2012, when it was, in addition, reviewed in detail in another paper of the same issue in the same journal.³ Besides, the system has also been addressed in another work that has tried to show that its real potential is actually great and to account for why, despite that

¹ See Peter Øhrstrøm, Jörg Zeller, and Ulrik Sandborg-Petersen, "Prior's Defence of Hintikka's Theorem. A Discussion of Prior's 'The Logic of Obligation and the Obligations of the Logician'," *Synthese* 188, 3 (2012): 449-454.

² Arthur N. Prior, "The Logic of Obligation and the Obligations of the Logician," *Synthese* 188, 3 (2012): 423-448.

³ Øhrstrøm et al., "Prior's Defence of Hintikka's Theorem," 449-454.

potential, people tend to reject its main consequence: Hintikka's theorem.⁴ The explanation about that tendency to rejection in this last work was based upon the mental models theory,⁵ a psychological theory basically stating that human reasoning is not logical and that essentially works taking semantic (in the linguistic sense of this word) possibilities into account. Nevertheless, that explanation concluded with a proposal to make people really understand Hintikka's theorem, which consisted of offering clear definitions for concepts such as 'permitted' or 'forbidden.' In this way, the chief idea was that concepts such as those ones were interpreted as referring to an impossibility to do a particular action in practice, that is, to a physical impossibility.

Nonetheless, the aim of this paper is only to give further support to the theorem and the demonstration of it that can be derived in Prior's system. In particular, it tries to reveal that, from other frameworks or methods, that theorem and that demonstration could also appear to be correct and acceptable. This will be done by means of the analysis of a concrete axiom in the system given by Prior and resorting to a particular approach of contemporary philosophy. The axiom is the one that can be deemed as the most relevant axiom in Prior's system, since, while obviously all of the elements in that system are important, it is the axiom that explicitly links the concept of necessity (and therefore a machinery such as the one that modal logics can provide today) to obligation (and therefore to a number of logical resources such as the ones that deontic logic can give nowadays). On the other hand, the approach is the well-known method of extension and intension proposed by Carnap.⁶

Of course, it is clear that a similar theoretical task could be done by considering other axioms (or even other elements) in Prior's system and other methods of analysis of meanings. However, the study that will be carried out here

⁴ Miguel López-Astorga, "What Is Possible and What Is permitted: Hintikka and Prior," *Analele Universitatii din Craiova, Seria Filosofie*, 39, 1 (2017): 57-66.

⁵ See, e.g., Sangeet Khemlani, Thomas Hinterecker, and Philip N. Johnson-Laird, "The Provenance of Modal Inference," in *Proceedings of the 39th Annual Conference of the Cognitive Science Society*, eds. Glenn Gunzelmann, Andrew Howes, Thora Tenbrink, and Eddy J. Davelaar (Austin: Cognitive Science Society, 2017), 259-264; Ana Cristina Quelhas and Philip N. Johnson-Laird, "The Modulation of Disjunctive Assertions," *The Quarterly Journal of Experimental Psychology* 70, 4 (2017): 703-717; Ana Cristina Quelhas, Célia Rasga, and Philip N. Johnson-Laird, "A Priori True and False Conditionals," *Cognitive Science* 41, 55 (2017): 1003-1030.

⁶ Rudolf Carnap, *Meaning and Necessity: A Study in Semantics and Modal Logic* (Chicago: The University of Chicago Press, 1947).

by using the mentioned axiom and framework can be illustrative enough to lead to think that more developments in this way, although they can always be appropriate and interesting, would be trivial and superfluous. In any case, to do what has been said, firstly the axiom will be described and, then, given that, as shown below, it has a conditional formal structure, two options will be taken into account. One of them will be the hypothetic scenario in which its antecedent is true, and the other one will be the possible alternative in which it is false. Evidently, each of these two options will be reviewed paying attention to the consequences that can be drawn from the theses and definitions included in Carnap's method and which can be related to what is provided by the axiom, the goal being, as stated, to show that that axiom is compatible with and can be assumed under a method such as that of Carnap. But the next section begins with the first action to do, which is, as also indicated, to explain what the axiom is exactly.

The Link between Necessity and Obligation in the System Proposed by Prior⁷

The particular axiom that will be addressed here is maybe the most important element in Prior's system, since it is, as pointed out, the element that relates modal logic to deontic logic in that system, and, accordingly, the element that, after all, allows relating what is possible (or impossible) to what is permitted (or unpermitted), and, in this way, the demonstration of Hintikka's theorem (which is not reproduced here again because, obviously, it is to be found in texts such as some of the ones that have been cited). It is as follows:

$$[I] \Box(p \rightarrow q) \rightarrow (Op \rightarrow Oq)$$

(Although with a different symbol for the conditional, [I] is axiom (3) in the text by Øhrstrøm et al.⁸).

The system is based on classical propositional calculus and hence ' \rightarrow ' in [I] stands for the material implication in that calculus, or, if preferred, in frameworks more or less akin to that presented by Deaño⁹. Nevertheless, perhaps the other symbols are more relevant for the aims of this paper. ' \Box ' is, as usual in modal logic, the operator of necessity, and, as also customary in that logic, it can be defined by the concept of possibility, whose symbol is ' \Diamond ':

⁷ Prior, "The Logic of Obligation," 423-448.

⁸ Øhrstrøm et al., "Prior's Defence of Hintikka's Theorem," 449-454.

⁹ Alfredo Deaño, *Introducción a la lógica formal* (Madrid: Alianza Editorial, 1999).

$$[II] \Box(x) =_{df} \neg\Diamond\neg(x)$$

(Where, obviously, ‘ \neg ’ represents negation and the ‘ x ’ between brackets to any well-formed formula in classical propositional logic).

And, as it is well known, possibility can also be defined by virtue of necessity in modal logic:

$$[III] \Diamond(x) =_{df} \neg\Box\neg(x)$$

As far as ‘O’ is concerned, clearly, it is the symbol of obligation in deontic logic, and, in this last logic, ‘P’, that is, the symbol standing for permission, habitually defines it:

$$[IV] O(x) =_{df} \neg P\neg(x)$$

But, as in the previous case, ‘P’ can be defined using ‘O’ too:

$$[V] P(x) =_{df} \neg O\neg(x)$$

Thus, it is absolutely clear the sense of [I]. It provides that the fact that a conditional is necessary implies that, if its antecedent is obligatory, then its consequent is obligatory as well. A formula such as that is not really hard to accept, and that can be seen by means of a method such as the one of Carnap.

L-truth and the Case in Which the Antecedent is L-true

A very important aspect of Carnap’s method is that it includes L-concepts. Those are concepts that, in a similar way as Kantian analytical judgments, are correct a priori and just by virtue of their meanings in the particular language that is being used. So, given a number of ‘state-descriptions’ (which is the expression to which Carnap resorts to indicate something similar to what the possible worlds are in modal logic), a L-concept is a concept that is correct in all of the state-descriptions. In this way, it can be said, for example, that (x) is L-true if and only if (x) is true in all of the state-descriptions that can be thought (see, e.g., Definition 2.2 in the text by Carnap¹⁰), and this leads one to note that what $\Box(x)$ actually provides is that (x) is L-true (see, e.g., Convention 39-3 in the text by Carnap¹¹).

Nonetheless, what is interesting now is that just a few notions such as these ones coming from the method of extension and intension can be sufficient to show that [I] should be accepted. Two possibilities can be thought in this regard: that its

¹⁰ Carnap, *Meaning and Necessity*.

¹¹ Carnap, *Meaning and Necessity*.

antecedent, $\Box(p \rightarrow q)$, is true and that it is false. The first case will be dealt with in this section and the second one in the next section.

If a formula such as $\Box(p \rightarrow q)$ is true, that means that $p \rightarrow q$ is true in all of the state-descriptions and that hence a state-description in which it is false cannot be thought ($p \rightarrow q$ is L-true). However, in classical logic, implication, as said, is material, and, as it is well known, that in turn means that,

$$[VI] (x) \rightarrow (y) =_{df} \neg[(x) \wedge \neg(y)]$$

(Where ‘ \wedge ’ stands for conjunction).

Thus, if the antecedent of [I], $\Box(p \rightarrow q)$, is true, by [VI], its consequent, $Op \rightarrow Oq$, also has to be true. In fact, it has to be true in exactly the same cases as the antecedent. Accordingly, since the antecedent is L-true, the consequent needs to be so too, and, therefore, $Op \rightarrow Oq$ has to be true in all of the state-descriptions as well. And this leads to another important concept in Carnap’s framework, which is the concept of L-equivalence: two formulae are L-equivalent when the state-descriptions in which they are true are exactly the same (see, e.g., Result 2-6 in the text by Carnap¹²). So, it can be said that this formula is L-true:

$$[VII] \Box(p \rightarrow q) \leftrightarrow (Op \rightarrow Oq)$$

(Where ‘ \leftrightarrow ’ is the symbol of the biconditional relationship).

But, following classical propositional calculus, this formula can also be drawn from [VII] and hence is L-true too:

$$[VIII] (Op \rightarrow Oq) \rightarrow \Box(p \rightarrow q)$$

and, by [IV], this is another L-true formula:

$$[IX] (\neg P\neg p \rightarrow \neg P\neg q) \rightarrow \Box(p \rightarrow q)$$

And, by propositional calculus, the same can be claimed for this sentence:

$$[X] (P\neg q \rightarrow P\neg p) \rightarrow \Box(p \rightarrow q)$$

Even, resorting to an equivalence that is also used in Prior’s system to derive Hintikka’s theorem (the equivalence, by [II], between $\Box(p \rightarrow q)$ and $\neg\Diamond(p \wedge \neg q)$; see Theorem (4) in the text by Øhrstrøm et al.¹³), it can be stated that the following is a L-true formula as well:

$$[XI] (P\neg q \rightarrow P\neg p) \rightarrow \neg\Diamond(p \wedge \neg q)$$

¹² Carnap, *Meaning and Necessity*.

¹³ Øhrstrøm et al., “Prior’s Defence of Hintikka’s Theorem,” 449-454.

Of course, more combinations are possible based upon the definitions indicated above and classical propositional calculus. However, those that have been pointed out can be enough to show the point of this section. An example with thematic content of a case in which $\Box(p \rightarrow q)$ is true, and hence L-true, has to be an example providing a meaning relationship between p and q . In this way, a simple material conditional would not be enough. Maybe it should be a conditional fulfilling criteria such as the one of Chrysippus of Soli¹⁴ or, more recently, the one of the strict implication proposed by Lewis.¹⁵ As it can be noted in his book, Carnap¹⁶ does not seem to avoid discussions in this direction. Nevertheless, perhaps what is important now is to highlight that, if $\Box(p \rightarrow q)$ is correct, as said, p and q cannot have any content. Their content has to be such that the combination p and $\neg q$ is impossible (in any state-description that can be thought).

Undoubtedly, different examples with thematic content of $\Box(p \rightarrow q)$ being true can be raised. Nonetheless, what is truly interesting here is that any of those examples appears to make sense and to be absolutely coherent not only with [I], but also with formulae [VII] to [XI]. One of those examples can be sufficient to see that. Given a sentence such as this one:

[XII] If I drink rum, then I drink alcohol.

Clearly, it is not possible that the antecedent is true and the consequent is false, as drinking rum necessarily implies drinking alcohol. And these contents for p and q do not seem to cause great difficulties to a formula such as [I], since what this last formula would provide would be that,

[I] The fact that it is necessary that, if I drink rum, then I drink alcohol implies that, if it were obligatory to drink rum, then it would be obligatory to drink alcohol.

¹⁴ E.g., Jonathan Barnes, Susanne Bobzien, and Mario Mignucci, "Logic," in *The Cambridge History of Hellenistic Philosophy*, eds. Keimpe Algra, Jonathan Barnes, Jaap Mansfeld, and Malcolm Schofield (Cambridge: Cambridge University Press, 2008), 77-225; William Kneale and Martha Kneale, *The Development of Logic* (Oxford: Clarendon, 1962); Robert R. O'Toole and Raymond E. Jennings, "The Megarians and the Stoics," in *Handbook of the History of Logic, Volume 1. Greek, Indian and Arabic Logic*, eds. Dov M. Gabbay and John Woods (Amsterdam: Elsevier, 2004), 397-522.

¹⁵ Clarence Irving Lewis, *A Survey of Symbolic logic* (Berkeley: University of California Press, 1918).

¹⁶ Carnap, *Meaning and Necessity*.

An Axiom Linking Necessity and Obligation Provided by Prior and Its Analysis...

Obviously, it is very hard to think about a world in which drinking rum is mandatory. However, the example appears to be totally coherent and provides an idea that is unquestionably correct: if [XII] is L-true, or, if preferred, if \Box [XII] is so, then, as indicated, in the hypothetical case in which its antecedent were obligatory, its consequent would be obligatory too. But something similar can be said with regard to [VII]:

[VII] The fact that it is necessary that, if I drink rum, then I drink alcohol is equivalent (or L-equivalent) to the fact that, if it were obligatory to drink rum, then it would be obligatory to drink alcohol.

Indeed, it is also difficult to imagine a state-description in which, [XII] being true (and, as pointed out, what this example of [VII] clearly states is that it is so in all the state-descriptions), it is not true, at the same time, that the obligation to drink rum would imply the obligation to drink alcohol, and vice versa. In this way, the case with [VIII] would not be very different:

[VIII] If it is true that, if it is obligatory to drink rum, then it is obligatory to drink alcohol, then it is also necessarily true that, if rum is drunk, then alcohol is drunk.

As claimed for [VII], it would not be an easy task to think about a circumstance in which this instance of [VIII] did not hold, which makes the case of [IX] obvious as well:

[IX] If it is true that, if it is not permitted not to drink rum, then it is not permitted not to drink alcohol, then it is also necessarily true that, if rum is drunk, then alcohol is drunk.

Maybe any comment on this last example would be trivial, since it is clear that its meaning and sense are not very different from those of the example given for [VIII]. And exactly the same can be stated in connection to [X], whose instance would be:

[X] If it is true that, if it is permitted not to drink alcohol, then it is permitted not to drink rum, then it is also necessarily true that, if rum is drunk, then alcohol is drunk.

And, finally, the application of the content of [XII] to [XI] also leads to a situation so similar to the previous ones as to make any explanation about it superfluous:

[XI] If it is true that, if it is permitted not to drink alcohol, then it is permitted not to drink rum, then it is also necessarily true that it is not possible to drink rum

and not to drink alcohol.

So, beyond the mentioned difficulty that to think about a hypothetical situation in which drinking alcohol is mandatory can raise, the examples with thematic content above allow checking that Carnap's semantic method to analyze meanings not only enables to accept the key axiom in Prior's system (and hence, as that fact seems to imply, his demonstration of Hintikka's theorem), but also to consider it to be absolutely suitable. As shown below, this does not greatly change if it is supposed that the antecedent of [I], $\Box(p \rightarrow q)$, is false.

L-truth and the Case in Which the Antecedent is L-false

But, even in the case that $\Box(p \rightarrow q)$ were L-false, [I] would keep being L-true. And the concept used now is 'L-false' because, according to Carnap,¹⁷ a formula is L-false if its negation is L-true, and it is obvious that, if $p \rightarrow q$ is not L-true, $\neg\Box(p \rightarrow q)$ is L-true, and hence $\Box(p \rightarrow q)$ is L-false.¹⁸

Certainly, as it is well known, in classical propositional calculus, in a consistent way with [VI], $(x) \rightarrow (y)$ is always true when (x) is not. So, if $\Box(p \rightarrow q)$ were untrue, [I] would be, in any state-description, true, no matter what the truth-value of $Op \rightarrow Oq$ were. Accordingly, the problem could be only in the cases [VIII] to [XI], in which $\Box(p \rightarrow q)$ —or, in [XI], the L-equivalent formula $\neg\Diamond(p \wedge \neg q)$ —is the consequent and, therefore, one might think, precisely by [VI], that the possibility exists that the formula in its entirety is false. It would be sufficient that the antecedent were true, since, thus, the antecedent would be true and the consequent would be untrue. However, situations such as this last one would not be really possible for at least two reasons.

Firstly, to modify the order of the clauses in a formula such as [I], it has to be transformed into [VII], and this is only possible if both the antecedent and the consequent are true in the same state-descriptions. So, if, for example, in [VIII], $Op \rightarrow Oq$ were true and $\Box(p \rightarrow q)$ were false, that would mean that they are not true in the same state-descriptions, that they are not equivalent (or L-equivalent), and that, therefore, neither [I] could be transformed into [VII] nor [VIII] could be derived from [VII]. And, of course, arguments very akin to these ones apply to formulae [IX] to [XI]. Obviously, the only possibility in which $\Box(p \rightarrow q)$ can be false and equivalent to $Op \rightarrow Oq$ is the case in which $Op \rightarrow Oq$ is false in the same

¹⁷ Carnap, *Meaning and Necessity*, e.g., Definition 2-3-a.

¹⁸ See also, e.g., Carnap, *Meaning and Necessity*, Convention 39-3.

state-descriptions as $\Box(p \rightarrow q)$, that is, in all of the state-descriptions. In this case, antecedent and consequent could change their places in the formula, but, given that, as said, the implication is material in classical propositional calculus, formulae [I] and [VII] to [XI] would continue to be true in all of the possible state-descriptions.

Secondly, if $\Box(p \rightarrow q)$ were L-false, then there would not be semantic relation (whether in the sense indicated by Chrysippus or the one indicated by Lewis or any other) between p and q . In this way, if $p \rightarrow q$ can be false, that is only because the antecedent and the consequent can be linked randomly. And that is what happens in sentences such as this one:

[XIII] If I wear hat, then I wear black shoes.

Clearly, the antecedent and the consequent are not semantically related in [XIII]. Hence they are not so either in [XIV]:

[XIV] If it is obligatory to wear hat, then it is obligatory to wear black shoes.

It is evident that neither [XIII] nor [XIV] can be L-true and that, accordingly, there can be state-descriptions in which one of them is true and the other one is false. Therefore, again, it is not possible to speak about equivalence (or L-equivalence) in the first place, it is not possible to transform [I] into [VII], it is not possible to draw formulae such as [VIII] to [XI] from [VII], and it is not possible that [I] leads to formulae with a true antecedent and a false consequent.

Conclusions

So, this paper can be deemed as one more piece of evidence showing that theorems such as the one of Hintikka may not be absurd and that Prior's system makes sense and has an interesting potential to be used. From a framework, in principle, different from the one of Prior,¹⁹ that of the method of extension and intension provided by Carnap, which, while it takes modality into account, does not consider deontic logic, his axiom seems to be admissible with complete justification. And this, regardless of the fact that it is also, as said, further support for the theorem raised by Hintikka, has two clear consequences.

On the one hand, it appears that Prior's system deserves to continue to be developed. Indeed, it seems to be correct from different perspectives, for example,

¹⁹ Prior, "The Logic of Obligation," 423-448.

from the one adopted by López-Astorga²⁰ and the one based upon Carnap's semantic method assumed here. Therefore, one might think that maybe, by working under its approach, interesting and surprising conclusions of great relevance in logic, philosophy, and science could be achieved.

On the other hand, this paper also appears to show that a lot of work remains to be done in fields such as linguistics and philosophy of language. As indicated above, the text by López-Astorga²¹ tries to clarify what certain words involved in the theorem presented by Hintikka and the system raised by Prior actually mean. That is the case of, for example, 'permitted', which is linked in López-Astorga's paper to senses such as impossibility to do an action from an ontological point of view. Thus, it can be thought that the problems of approaches such as those of Hintikka and Prior are that they refer to several different levels, including physics, metaphysics, ethics, and linguistics, and that perhaps only interdisciplinary studies paying attention to most of those levels can truly reveal all the richness that those approaches have and the real meanings of the words, concepts, and operators used in them, which do not denote exactly the same in all of such levels. In any case, maybe another fact that confirms that the difficulties with these issues are related to the need for a clarification of what certain elements really mean in systems such as the one of Prior is that the method used in this paper is a method to study in-depth and recover the meanings of expressions. So, it is possible that, beyond the method one follows, this is the chief task to do in the near future.²²

²⁰ López-Astorga, "What Is Possible and What Is Permitted," 57-66.

²¹ López-Astorga, "What Is Possible and What Is Permitted," 57-66.

²² This paper is a result of the Project CONICYT/FONDECYT/REGULAR/FOLIO N° 1180013, "Recuperación de las formas lógicas de los enunciados a partir de un análisis de las posibilidades semánticas a las que hacen referencia", supported by FONDECYT (National Fund for Scientific and Technological Development), Government of Chile.

SCIENCE, VALUES, AND THE PRIORITY OF EVIDENCE

P.D. MAGNUS

ABSTRACT: It is now commonly held that values play a role in scientific judgment, but many arguments for that conclusion are limited. First, many arguments do not show that values are, strictly speaking, indispensable. The role of values could in principle be filled by a random or arbitrary decision. Second, many arguments concern scientific theories and concepts which have obvious practical consequences, thus suggesting or at least leaving open the possibility that abstruse sciences without such a connection could be value-free. Third, many arguments concern the role values play in inferring from evidence, thus taking evidence as given. This paper argues that these limitations do not hold in general. There are values involved in every scientific judgment. They cannot even conceivably be replaced by a coin toss, they arise as much for exotic as for practical sciences, and they are at issue as much for observation as for explicit inference.

KEYWORDS: values and science, inductive risk, ampliative risk, epistemic values, evidence, externality

Introduction

Recent philosophical literature on science and values has shown numerous ways in which science is not and could not be value-free.¹ The point is not just that scientific judgment respects the value of evidence and theoretical virtues, but also that it reflects whether certain outcomes would be good or bad. Nevertheless, arguments for this often yield only a limited conclusion.

First, some arguments appeal to values as a means of choosing theories in the face of empirical underdetermination. Yet, as critics note, it does not need to be values that fill the gap. In principle, a scientist faced with underdetermination between two options might instead just flip a coin. Call this the *randomizer reply*.

Second, arguments for the value-laden nature of science often apply to some but not all science. For example, Anna Alexandrova argues that sciences of well-

¹ Elliott provides a systematic discussion of different ways that science and values may be connected; see Kevin C. Elliott, *A Tapestry of Values: An Introduction to Values in Science* (Oxford: Oxford University Press, 2017).

being (ones which make claims about health, for example) necessarily and legitimately reflect normative judgments.² Hilary Putnam similarly argues that sciences which employ thick concepts undercut the distinction between fact and value.³ These arguments do not apply, nor are they meant to apply, to abstruse sciences like astronomy or particle physics. The concept of the neutrino, for example, does not seem to have any intrinsically normative dimension. Call this the *policy-relevance restriction*.

Third, many arguments show only that values enter into inferences from evidence to an underdetermined conclusion. The evidence itself is taken as given. The arguments, Kevin Elliott and Daniel McKaughn note, "incorporate nonepistemic values only as a secondary consideration for resolving epistemic uncertainty."⁴ Matthew Brown decries these arguments for bringing in values too late. Adopting Brown's phrase, call this limitation the *lexical priority of evidence* over values.⁵

In sections 2 and 3, I consider several arguments that values necessarily enter into scientific inference. These arguments, as they are usually posed, treat evidence and values as separate inputs to the process. In section 4, I consider the role of values in scientific observation itself. Heather Douglas provides a clear example in which scientists had to judge whether prepared slides showed tumors or not.⁶ The judgments were significant for environmental regulation, and so values were involved. Douglas argues that similar entanglement would not obtain in science without policy implications. In sections 5-7, I argue that the connection between scientific judgment and values which is typified in ampliative inference and Douglas' example holds for all scientific observation. Scientists always have a choice about how to state an observation, between more significant but riskier formulations and less significant but safer ones. This choice always involves weighing values in the sense of the goodness or badness of various possible

² Anna Alexandrova, "Can the Science of Well-Being Be Objective?" *The British Journal for the Philosophy of Science* 69, 2 (2018): 421-445.

³ Hilary Putnam, *The Collapse of the Fact/Value Dichotomy and Other Essays* (Cambridge, Massachusetts: Harvard University Press, 2004).

⁴ Kevin C. Elliott and Daniel J. McKaughn, "Nonepistemic Values and the Multiple Goals of Science," *Philosophy of Science* 81, 1 (2014): 2.

⁵ Matthew J. Brown, "Values in Science Beyond Underdetermination and Inductive Risk," *Philosophy of Science* 80, 5 (2013): 829-839.

⁶ Heather E. Douglas, "Inductive Risk and Values in Science," *Philosophy of Science* 67, 4 (2000): 559-579.

outcomes. This weighing could not be done without values (contra the randomizer-reply), it occurs in all science (contra the policy-relevance restriction), and it occurs in the very formulation of the evidence (contra the lexical priority of evidence).

Underdetermination and Tie-Breaking

Some arguments connect science and values in cases where the usual standards of evidence are insufficient to decide among competing hypotheses. The arguments hold that, in such cases, scientists may responsibly select the hypothesis which best accords with their value commitments. Call this the *tie-breaker argument*. It applies only to cases where hypotheses score equally well with respect to the evidence, and values enter only to break the tie.⁷

A more subtle version of the tie-breaker argument is given by Helen Longino.⁸ She notes that connecting scientific theories to observable phenomena almost always requires auxiliary hypotheses, an aspect of underdetermination sometimes called the Duhem-Quine Problem. Values enter by way of value-laden auxiliary hypotheses, when matters cannot be settled by observation and value-neutral auxiliary hypotheses. Like the less-subtle tie-breaker argument, values are in play only when evidence does not determine theory choice.

Let's consider two standard replies to the tie-breaker argument.

One response is to claim that scientists could break ties by using a randomizer instead of making value judgments. This response has a long history, although it seems to have been suggested independently by different thinkers. Otto

⁷ Magnus and Longino situate these arguments in a broader analysis of 'underdetermination;' see P.D. Magnus, "Underdetermination and the Claims of Science" (PhD diss., University of California, San Diego, Department of Philosophy, 2003) and Helen Longino, "Underdetermination: A Dirty Little Secret," *STS Occasional Papers* 4 (2016), Department of Science and Technology Studies, University College London.

⁸ Helen Longino, *Science as Social Knowledge* (Princeton, NJ: Princeton University Press, 1990) and "Underdetermination."

Neurath suggests drawing lots.⁹ Gregor Betz suggests rolling a die.¹⁰ Inmaculada de Melo-Martín and Kristen Intemann suggest flipping a coin.¹¹

The idea of this randomizer reply is that underdetermination only shows that something besides evidence alone must determine theory choice. Although we could follow our value commitments to select preferable theories or auxiliary hypotheses, we could instead use a procedure that is independent of our values. Randomly selecting one theory over its competitors (or one set of auxiliary hypotheses over alternate sets) makes the choice without regard to which would be better or which we would prefer.

Nevertheless, adopting such a policy would be a practical decision. As an analogy, consider a mundane case in which I cannot decide which of two restaurants to visit for lunch and so flip a coin. My values and practical reasons have no influence over the outcome of the coin toss, of course, and so the selection is value-free *to that extent*. However, my values and preferences are involved in my decision to use coin-flipping as a way of resolving the choice. I want to go to lunch at one of two places, and I do not want to spend too much time or energy deciding. If someone asks why I went to one restaurant rather than the other, a complete answer would refer not just to the random process but also to the reasons I had for adopting that method. Similarly, deciding to flip a coin in the face of underdetermination would be practical and value-driven. Breaking ties by flipping coins would keep values from directly deciding specific winning hypotheses, but it would not ultimately escape the intrusion of values and practical decisions into theory choice.¹² Just as values might lead us to prefer a specific outcome or some auxiliary hypotheses, values might lead us to choose a random method. Choosing to believe *the outcome determined by the coin toss* when evidence itself underdetermines theory choice is still a value-driven decision. So the randomizer reply fails.

⁹ Otto Neurath, "The Lost Wanderers of Descartes and the Auxiliary Motive: On the Psychology of Decision," in *Philosophical Papers: 1913-1946*, eds. Robert S. Cohen and Marie Neurath (Dordrecht: D. Reidel, 1983 [1913]), 1-12.

¹⁰ Gregor Betz, "In Defence of the Value Free Ideal," *European Journal for the Philosophy of Science* 3, 2 (2013): 210.

¹¹ Inmaculada de Melo-Martín and Kristen Intemann, "The Risk of Using Inductive Risk to Challenge the Value-Free Ideal," *Philosophy of Science* 83, 4 (2016): 505.

¹² This is one way to read Neurath's argument: Neurath thinks that we would prefer an epistemic culture that does not let our preferences directly decide theory choice. So, he argues, we should draw lots. Preferences are still at work, at the level of general policy.

A second response to the tie-breaker argument is to insist that scientists should never break ties with anything but further evidence. Rather, scientists should be agnostic when evidence is insufficient to decide between rival hypotheses. They should pursue each hypothesis until evidence is uncovered which breaks the tie without any appeal to values. Call this the *wait-and-see reply*.

In answer to the wait-and-see reply, note that remaining agnostic and collecting further evidence is not always possible. As a practical matter, we cannot pursue every hypothesis. Pursuing hypotheses that figure in separate research programs may require different techniques, and so different investments in training and equipment. Although "thought experiments can be risked without hesitation," Neurath notes, it is not possible "in the same way, to train for more than one career."¹³ We could organize the scientific community so that different scientists were trained to pursue different research programs, but there is a limit to how many scientists can be trained and how many laboratories can be outfitted. Moreover, waiting for more evidence takes time. Brown argues that there are cases in which "we cannot wait for the end of inquiry for scientists to accept or reject a hypothesis, we cannot depend on anyone else to do it, and we must contend with uncertainty and underdetermination [... S]cientists find themselves in the business of accepting and rejecting hypotheses in such conditions."¹⁴ Even if we have the option to wait for compelling evidence, the cost of waiting might be higher than the probable cost of using quicker but less-reliable methods.¹⁵

One may still insist that it is appropriate (insofar as possible) to respond to underdetermination by remaining agnostic and collecting more evidence. This would limit the scope of the tie-breaker argument to cases where agnosticism is impractical. It would still be a regulative ideal to abide with unbroken ties until further evidence could be uncovered to change the score.

Retreating to agnosticism in this way would avoid using values to break ties between rival theories, but it would require that the evidence itself be value-free. Brown observes that the tie-breaker argument "begin[s] from a situation where the evidence is fixed and take[s] values to play a role in the space that is left over."¹⁶ The underdetermination of theory by data, in its very formulation, concerns the inference to theory once data are given. It presumes that evidence comes first,

¹³ Neurath, "The Lost Wanderers of Descartes," 3.

¹⁴ Brown, "Values in Science," 831-2.

¹⁵ Elliott and McKaughn, "Nonepistemic Values."

¹⁶ Brown, "Values in Science," 834.

before the question arises of whether values should be involved. Brown calls this presumption the *lexical priority of evidence over values*. He explains, "lexical priority means that evidence will always trump values."¹⁷ The picture is one on which evidence and values might independently underwrite our preferring one theory over another, and the lexical priority of evidence means that considerations of evidence always override considerations of values. The preferences underwritten by values only hold sway when evidence is silent and we are forced to make a choice.

The wait-and-see reply works to limit the conclusion of the tie-breaker argument in two respects: First, values are allowed to play a role only for questions of practical importance but not for abstruse matters (the policy-relevance restriction). Second, values play a role only after the evidence itself is given (the lexical priority of evidence).

Ampliative Risk and Our Duties as Knowers

In drawing an inference from evidence, there is inevitably a tension between striving to believe the truth and striving to avoid error. Scientists might be quick to judge or more cautious. If they are too quick to judge, they risk believing in error; such a result is a *false positive* or *type I* error. If they remain agnostic, they risk the opportunity cost of not having an accurate belief that they could have had; such a result is a *false negative* or *type II* error. This tension is resolved only by assessing what the cost would be of each possible error—that is, by reckoning with values. Therefore, values enter into scientific inference. Call this the *ampliative risk argument*.¹⁸

The impetus to believe truth drives us to adopt the claim best supported by evidence, but the impetus to avoid error drives us to wait and demand more evidence. William James puts it in histrionic terms: "Believe truth! Shun error!—these, we see, are two materially different laws; and by choosing between them we

¹⁷ Brown, "Values in science," 836.

¹⁸ It is standard, following Hempel, to call this the argument from inductive risk. I've opted for the label 'ampliative risk' because 'induction' is ambiguous between a narrow use (enumerative induction) and a broad use (ampliative inference) (Carl G. Hempel, "Science and Human Values," in *Aspects of Scientific Explanation and other Essays in the Philosophy of Science* (New York: The Free Press, 1965), 92). See also Kevin C. Elliott and Ted Richards, *Exploring Inductive Risk* (Oxford: Oxford University Press, 2017).

may end up coloring differently our whole intellectual life."¹⁹ A similar point is made by Richard Rudner, who writes that "our decision regarding the evidence... is going to be a function of the *importance*, in the typically ethical sense, of making a mistake in accepting or rejecting the hypothesis."²⁰ Heather Douglas puts the point in less grandiose terms: "Within the parameters of available resources and methods, some choices must be made, and that choice should weigh the costs of false positives versus false negatives. Weighing these costs legitimately involves social, ethical, and cognitive values."²¹

The tie-breaker argument was (at least partially) defused because not all ties need to be broken. We could wait and see. The argument here turns on the fact that waiting and seeing would itself be a choice, with benefits but also costs for our overall system of belief. As a result, values are always involved in assessing "the sufficiency of evidence, the weighing of uncertainty, and the consequences of error."²² Douglas calls this an *indirect role* for values.

Elsewhere, I have called this the James-Rudner-Douglas or JRD thesis: "Anytime a scientist announces a judgment of fact, they are making a tradeoff between the risk of different kinds of error. This balancing act depends on the costs of each kind of error, so scientific judgment involves assessments of the value of different outcomes."²³ The JRD thesis underwrites the ampliative risk argument but, as I will argue below, it is broader and more fundamental.

Note that one could not get around the JRD thesis by flipping a coin in cases where the evidence is equivocal. Every ampliative inference involves reckoning with the risks of different kinds of error. Errors are always logically possible, precisely because the inference is ampliative. It is possible for all of the premises to be true, but for the conclusion still to be false. So one would be flipping coins always and for everything.

¹⁹ William James, "The Will to Believe," in *Essays in Pragmatism*, ed. Alburey Castell (New York: Hafner Publishing Co., 1948), 100.

²⁰ Richard Rudner, "The Scientist qua Scientist Makes Value Judgments," *Philosophy of Science* 20, 1 (1953): 2.

²¹ Heather E. Douglas, *Science, Policy, and the Value-free Ideal* (Pittsburgh: University of Pittsburgh Press, 2009), 104.

²² Douglas, *Science, Policy*, 103.

²³ P.D. Magnus, "What Scientists Know Is Not a Function of What Scientists Know," *Philosophy of Science* 80, 5 (2013): 845. See also P.D. Magnus, "Science and Rationality for One and All," *Ergo* 1, 5 (2014): 129-138.

The ampliative risk argument as it is typically posed presumes that scientific inference leads to accepting or rejecting hypotheses.²⁴ The costs and benefits which scientists are supposed to reckon with are conditional on making the right or wrong choice. So a standard rebuttal is to deny that scientists should ever be flatly accepting or rejecting hypotheses.

Gregor Betz argues that scientists should report hedged hypotheses.²⁵ The idea is that, instead of reporting a categorical result 'H,' scientists should instead report something like 'Evidence strongly suggests but does not decisively confirm H.' Notice, however, that hedged claims still express a degree of confidence. Scientists must decide how much to hedge. Less hedging is riskier, but more hedging courts triviality. In order to escape all danger of being wrong, scientists might refuse to make any claims at all. As Stephen John observes, the price of that epistemic security would be "policy-impotence," having nothing to say that could be of any use to policy-makers or anyone else.²⁶ Scientists should, to use Betz's phrase, "simply admit their complete ignorance" in cases where they have no evidence whatsoever—but it would be pathological for them to plead complete ignorance just so as to avoid any chance of being wrong.²⁷

Responding to Rudner, Richard Jeffrey argues that scientists should only report probabilities.²⁸ Even though this allows scientists to avoid deciding for or against H, they must still decide for or against the claim that the probability of H is p . A scientist might hedge this report by returning an interval $p \pm e$ rather than a precise probability p . As the size of the interval is larger, the claim is safer but less useful in guiding action. The limit case, reporting that the probability of H is between 0 and 1, is a useless tautology. It does not even help to assume, as Bayesians sometimes do, that a scientist always has some degree of belief p in every hypothesis H. Since this could reflect their prior credence more than the weight of evidence, a scientist must still judge that it reflects enough evidence to merit reporting.

²⁴ For the argument formulated in terms of acceptance and rejection, see (e.g.) Rudner, "The Scientist qua Scientist;" Hempel, "Science and Human Values," 92-3; Brown, "Values in Science."

²⁵ Betz, "In Defence."

²⁶ Stephen John, "The Example of the IPCC Does Not Vindicate the Value Free Ideal: A Reply to Gregor Betz," *European Journal for the Philosophy of Science* 5, 1 (2015): 9.

²⁷ Betz, "In Defence," 9.

²⁸ Richard Jeffrey, "Valuation and Acceptance of Scientific Hypotheses," *Philosophy of Science* 23, 3 (1956): 237-246.

This reply is given already by Rudner, who argues that a hedged report is itself "nothing more than the acceptance by the scientist of the hypothesis that the degree of confidence is p or that the strength of evidence is such and such."²⁹ The question is precisely what the benefit would be of making a less-hedged, more-confident report (if it were true) and what the cost would be (if it were false). Settling on any particular conclusion, even if it is a conclusion about probabilities or the weight of evidence, is subject to ampliative risk. The JRD thesis applies, and so values play a role.

Brown objects that the ampliative risk argument (like the tie-breaker argument) presumes the lexical priority of evidence over values.³⁰ Note, however, that there can be no question here of evidence trumping values. The JRD thesis means that there can be no scientific conclusion without (at least implicitly) weighing the costs of various possible errors. There is no choice favored by the evidence alone, so *a fortiori* it makes no sense for that choice to presumptively win out.³¹

The argument from ampliative risk does not yield the blanket conclusion that, as Stijn Conix puts it, "it does not make sense to think of values and epistemic standards as taking priority over each other."³² It makes sense to think of some values that way. If we rank theories according to how well they promote a conception of human autonomy and according to how simple they are, for example, then these separate rankings might pick out different theories as best. The values could give separate preference orderings. The point of the JRD thesis is that there are *some* values which cannot be separated in this way. The enthusiasm to reach for a possibly true belief or the risk-aversion which makes one remain agnostic even as evidence accumulates—these values do not rank the possible

²⁹ Rudner, "The Scientist qua Scientist," 4, emphasis in original.

³⁰ Matthew J. Brown, "Values in Science." Brown calls the tie-breaker argument "the gap argument," and the ampliative risk argument is his "error argument."

³¹ ChoGlueck argues that the error argument (the argument from ampliative risk) is just a special case of the gap argument (the tie-breaker argument). This is directly rebutted by the fact that the tie-breaker argument is vulnerable to Brown's worry about the lexical priority of evidence in a way that the argument from ampliative risk is not. Christopher ChoGlueck, "The Error Is in the Gap: Synthesizing Accounts for Societal Values in Science," *Philosophy of Science* 85, 3 (2018): 704-725.

³² Stijn Conix, "Radical Pluralism, Ontological Underdetermination, and the Role of Values in Species Classification" (PhD diss., University of Cambridge, Queen's College, Department of History and Philosophy of Science, 2017), 102.

choices apart from epistemic standards. Enthusiasm may make a scientist respond to preliminary evidence with belief when their colleagues demand more evidence, but it is precisely a disagreement about how much evidence is *enough*. Risk-aversion will delay accepting a belief as evidence accumulates, but only an utter sceptic would refuse to believe regardless of how much evidence there might be. These values are entangled with the application of epistemic standards, so it makes no sense to think of one as taking priority over the other.

A different way to construe the complaint about lexical priority is that the evidence has been presumed to be value-free. Although evidence and the costs of possible errors both enter into the epistemic calculation, they do so as independent variables. One may complain, as Elliott and McKaughn do, that the JRD thesis adds "nonepistemic values only as a secondary consideration."³³

This objection is especially apt when the ampliative risk argument is posed in terms of type I and type II errors. These terms come from statistical hypothesis testing, where the problem is to specify a rule for accepting or rejecting hypotheses given a data set. The data set itself is not at issue. This construal of the argument is encouraged by Rudner's insistence that "every scientific inference is properly construable as a statistical inference."³⁴ It is also encouraged just by posing the argument in terms of *inductive risk*. 'Induction' and 'inductive inference' are often used narrowly to pick out inference from a sample to a population or from a finite track-record to a generalization. From given observations, the inductive problem is how to generalize or draw conclusions. Evidence is lexically and literally prior.

Of course, James and Douglas do not pose the argument in terms of statistical inference. Their arguments apply to ampliative inference generally, rather than just to inductive inference narrowly-construed. Nevertheless, posing the argument in terms of *inferential risk* makes it turn on the move from evidence to hypothesis. Douglas specifies that values, in an indirect role, "determine the importance of the inductive gaps left by the evidence."³⁵ Even though evidence does not recommend a conclusion without some values, the evidence is in a sense primary.

³³ Elliott and McKaughn, "Nonepistemic Values," 2.

³⁴ Rudner, "The Scientist qua Scientist," 3.

³⁵ Douglas, *Science, Policy*, 96.

Evidence That Matters for Policy

Douglas offers an example that rebuts this as a general worry. In the 1970s, slides of rat livers were prepared as part of a study of dioxin toxicity. These slides were evaluated by different teams of scientists over more than a decade, and different numbers of liver tumors were reported in the different evaluations. Douglas writes,

Although not as formal as setting a level for statistical significance, the pathologists must be similarly concerned with false positives and false negatives. Suppose a pathologist chooses to take all borderline cases and judge them to be non-cancerous lesions. ... The consequences for such an approach will be an underestimation of malignancies and thus an underestimation of risk.³⁶

These slides were revisited again and again precisely because the study was "important in regulation," so that an estimate of lower risk would "likely lead to a relaxed regulation" which could "cause increased harm to the public." Conversely, judging borderline cases as malignant would have erred on the side of protecting public health "at the economic costs of potentially unnecessary regulation."³⁷

In the spirit of Jeffrey and Betz, one might note that scientists could reject the requirement that slides be sorted decisively into those that showed acute toxicity and those that did not. However: Although scientists could emphasize their uncertainty and the tentativeness of their conclusions to different degrees, there was no neutral way of relaying the objective situation to policymakers. Whatever report scientists gave, even refusing to report at all, would have consequences. So hedging or reporting confidence intervals could not escape the practical significance of reporting their results in one way rather than another.

This is kind of an easy case, though. Scientists knew that their observations would have consequences for regulation and public health. Their observations had a clear valence in practical and ethical terms. Douglas herself notes this policy-relevance restriction in the scope of her argument. She writes that

there are some areas of science where making a wrong choice has no impact on anything outside of that area of research. One may think, for example, of research into the coherence properties of atom beams. It is very difficult to fathom how errors in such research could have non-epistemic consequences. Hence, scientists doing such research need not consider non-epistemic values.³⁸

³⁶ Douglas, "Inductive risk," 571.

³⁷ Douglas, "Inductive risk," 571.

³⁸ Douglas, "Inductive risk," 577.

Lots of science does not have any direct practical consequences or foreseeable application. Douglas claims that, in such cases, decisions can be made on strictly epistemic grounds. In the next section, I argue that this concedes too much. Even in exotic sciences like particle physics, the JRD thesis applies.

Observation and Externality

Following Trevor Pinch, we can distinguish possible observation reports by their degree of *externality*.³⁹ Externality is the inverse of immediacy—that is, an observation posed at a lower degree of externality is in more direct terms.

A high externality report is riskier but potentially more significant. A lower externality report, in contrast, is safer but less interesting. A scientist initially reports their observation at some level of externality. If their report is challenged, they can redescribe the observation at a lower level of externality and offer an argument which takes the lower-externality report as a premise and yields the higher-externality report as a conclusion.

Pinch gives the example of Ray Davis' work to detect solar neutrinos in the 1960s. The detection was a complicated operation. A large tank of tetrachloroethylene, stored in an abandoned mine shaft, was used as a target. Some of the chlorine atoms interacted with solar neutrinos to produce an isotope of argon (argon-37). The accumulated argon-37 was extracted from the tank and was measured based on its characteristic decay. Finally, there were some outputs from instruments—Pinch refers to them as "splodges."

The outcome of the work could be reported at different levels of externality. From higher to lower externality, the report might be a claim about:

- The rate of particular reactions in the sun
- Neutrinos generated in the sun
- Neutrinos arriving at Earth
- Argon-37 atoms in the tank
- Splodges on the apparatus

This list reflects the levels parsed out by Pinch, except that he groups the second and third together just as "Solar neutrinos."⁴⁰ The distinction between

³⁹ Trevor Pinch, "Towards an Analysis of Scientific Observation: The Externality and Evidential Significance of Observational Reports in Physics," *Social Studies of Science* 15 (1985): 3-36.

⁴⁰ Pinch, "Towards an Analysis," 17.

neutrinos generated in the sun and neutrinos arriving at Earth was ultimately important, however. Solar neutrino oscillation, a change in neutrinos as they travel to Earth from the sun, is now taken to explain Davis' result.⁴¹

We could also parse this more finely. If Davis reported the number of argon-37 atoms in the tank and someone challenged that report, for example, he could explain how it can be inferred from the number of argon-37 atoms extracted from the tank or from the radioactivity of the extracted material. This would report the observation at levels of externality between the last two in the list above.

And we might add further levels of even lower externality. Faced with scepticism about the external world, Davis might report his sense data and argue on that basis that there were splodges. At that extreme, he would no longer be reporting anything of scientific interest.

Davis set out to test theories about what was going on in the sun, so the scientifically most interesting claim would be an observation of specific reactions in the sun. This would involve considerable risk, however, because there were all sorts of ways in which his report about solar neutrinos could turn out to be wrong. This is not just in-principle scepticism, either, since he observed far fewer solar neutrinos than physicists predicted based on the reactions that they expected were happening. At the lowest degrees of externality, however, the report risks being trivial. The large and complicated project would hardly be justified if, at the end of it, he could report nothing more than splodges. So characterizing the observation required balancing considerations of risk against considerations of significance.

These considerations reflect the cost of believing a risky claim (if it were false) and the cost of forgoing a significant claim (if it were true). These are the simultaneous demands that we should avoid error and seek truth. The JRD-thesis, that this tension is inescapable, applies as much to the observation claim as to the conclusions of inference. This means that values enter not just into the inference from evidence, but into stating the evidence itself. In considering Jeffrey and Betz, above, we saw that scientists can assign lower probabilities or hedge their reports in order to trade significance for security. Moving to lower levels of externality is a distinct strategy for doing so. In this example, Davis could have made his report safer by widening the error bars on his observation report of neutrino flux or by maintaining precision while characterizing the observation as being about objects

⁴¹ Note that reports at different levels of externality may exhibit the same degree of generality. For example, the rate of neutrinos generated in the sun is just as specific and concrete as the rate of neutrinos arriving at Earth.

at lower levels of externality. Just as widening the error bars too far would make the report trivial, so too would retreating to the lowest levels of externality.

Scientists in the decades following Davis' observations worked to figure out what he had observed, both what could be concluded from it but also what the correct description of it was. Settling on the right level of externality took decades.⁴² Davis could not wait for these developments, but had to decide what to believe and what to report at the time. He faced uncertainty, and navigating it required weighing the potential costs and benefits of different possible beliefs.

To generalize, the argument is this: In stating their observations, scientists face a choice between different levels of externality. Observation reports posed at a high level of externality are more significant, so a scientist who declines to accept a report in those terms risks missing out on the chance to believe an important truth. Conversely, such reports are also riskier, so a scientist who accepts such a report risks believing something false. There is no purely epistemic rule which assigns costs to these risks. Instead, they are a matter of value judgment. Just as ampliative risk means that values always play a role in theory choice, values play a role in every observation.

This same pattern is seen in other cases of experimental science, and the concept of externality is helpful for describing what is going on in general terms.

Kent Staley discusses collaborative research by groups like the one at Fermilab which discovered key evidence for the top quark in the 1990s.⁴³ They must settle on reporting their findings in some form. This yields two conflicting but indispensable pressures, Staley argues: "[T]hey seek to avoid the embarrassment of making claims that subsequent work reveals to be false; they also seek to achieve prominence and esteem by making novel and significant claims that are upheld by further critical scrutiny."⁴⁴ Claiming to have observed the top quark was a high externality report, significant but also risky.

Boaz Miller notes that experimentalists must distinguish signal from noise.⁴⁵ There are different methods for reducing raw data, and selecting which method to

⁴² For a summary of subsequent developments, see John N. Bahcall, "Solving the Mystery of the Missing Neutrinos," April 28, 2004. http://www.nobelprize.org/nobel_prizes/themes/physics/bahcall/

⁴³ Kent W. Staley, "Evidential Collaborations: Epistemic and Pragmatic Considerations in 'Group Belief,'" *Social Epistemology* 21, 3 (2007): 321-335 and "Decisions, Decisions: Inductive Risk and the Higgs Boson," in *Exploring Inductive Risk*, eds. Elliot and Richards, 37-55.

⁴⁴ Staley, "Evidential Collaborations," 323.

⁴⁵ Boaz Miller, "Catching the Wave: The Weight-Adjusting Account of Values and Evidence,"

use is subject to social influences and considerations of risk. Claiming only to have observed the raw data is a low externality report, safe but insignificant. Claiming to have observed a definite signal is higher externality, potentially significant but also potentially wrong.

Objections on the Basis of Epistemic Values

In this section, I rebut some possible objections which suggest that the values involved in picking a level of externality are only epistemic values.

Here is a first try at such an objection: Describing epistemic values, Ernan McMullin writes, "One value, namely truth itself, has always been recognized as permeating science."⁴⁶ Avoiding error, too, is clearly an epistemic matter. So, one may object that I haven't shown how non-epistemic or ethical values are in play.

However, the epistemic duties to pursue truth and avoid error ultimately pull in opposite directions. So seeking truth and shunning error are not enough, by themselves, to determine belief. When they are in conflict, the epistemic duties themselves cannot tell us how to strike a balance, and we must consider how important it would be to believe a claim (if it were true) and what the cost would be of remaining agnostic (avoiding error, if the candidate belief were false). This *importance* and *cost* will be practical rather than narrowly epistemic. Selecting a level of externality requires weighing enthusiasm against caution—that is, it requires reckoning with values.

Staley describes the risk of embarrassment pitted against the thirst for esteem—practical considerations for scientists, rather than merely epistemic concerns.⁴⁷ To take a schematic example, a junior scientist who needs to publish in order to have a chance at tenure might favor reporting a result now rather than waiting for further proof. They might be wrong, but waiting could be tantamount to dooming their career. Or consider a somewhat different junior scientist who has published enough to secure tenure. The risk of publishing too soon and detracting from their overall CV might lead them to greater caution than the first scientist. In both cases, they would be acting on their duties to make true claims and avoid making false ones, but in neither case is it just truth and falsity that guide their

Studies in History and Philosophy of Science Part A, 47 (2014): 69–80.

⁴⁶ Ernan McMullin, "Values in Science," *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, Volume Two: Symposia and Invited Papers* (1982), 6.

⁴⁷ Staley, "Evidential Collaborations," 323, cited above.

decisions. They are weighing the expected utility of making claims or refraining from doing so. Therefore, the values involved go beyond the merely epistemic.

Note again that the argument does not show that every value or utility is relevant to scientific judgment.⁴⁸ Imagine that one of the junior scientists has already published a paper. It might be good for them if the claims that they made in the paper were true, but that by itself does not give them any reason to believe the claims. The relevant costs and benefits are the conditional ones: what they would have to gain by believing it *if it were true*, what would it cost them to believe *if it were false*, and so on.⁴⁹ McMullin is explicit that he would view scientists weighing the expected utilities of various judgments as an intrusion of ethical (non-epistemic) values.⁵⁰

To revise the objection: One might concede that seeking truth and shunning error alone are not enough to settle theory choice but appeal to a longer list of epistemic values.

It is hard to reply to this without some candidate for what these extra epistemic values could be. Standard lists of theoretical virtues include things like fit with evidence, coherence, consistency, simplicity, scope, and fertility.⁵¹ However, as theoretical virtues, these do not readily apply to observation. I do not see how these would provide *any* guidance in a case like Davis' neutrino observation, nonetheless enough guidance to settle the appropriate level of externality. One might attempt to enlarge the list of epistemic values even further to include other scientifically significant considerations, but I do not see how this could be made to work without collapsing the distinction between epistemic and non-epistemic values.⁵²

⁴⁸ My conclusion here is more modest than Staley's claim that "the decision about communicating the outcome of an experiment is subject to the full range of utility considerations applicable to any practical decision" (Staley, "Decisions, Decisions," 53).

⁴⁹ This difference between categorical and conditional values is Douglas' distinction between values in a direct and values in an indirect role.

⁵⁰ McMullin, "Values in Science," 8.

⁵¹ The list reflects ones given by Kuhn, McMullin, and (in a critical vein) Longino (Thomas S. Kuhn, "Objectivity, Value Judgment, and Theory Choice," in *The Essential Tension: Selected Studies in Scientific Tradition and Change* (Chicago: University of Chicago Press, 1977), 320-339; McMullin, "Values in Science;" Longino, *Science as Social Knowledge*). Longino and Douglas argue that there is no legitimate distinction between values like these and any others we might revere, but I accept the distinction for the sake of argument.

⁵² See Philip Kitcher, *Science, Truth, and Democracy* (Oxford: Oxford University Press, 2001),

To revise the objection again: One might insist that community standards settle the relevant considerations. There are norms in the scientific community for when and what to report, and the two junior scientists considered above only have so much leeway in what claims to publish. In the schematic case, there are professional standards which limit how enthusiastic the first scientist is allowed to be.

Although community standards might make individual values irrelevant in particular cases, the general standards themselves must strike some balance between different possibilities of error. We saw above that a practical policy of flipping a coin would still reflect values at the level of policy. The same holds, for example, for setting a statistical threshold that results must meet in order to be publishable.

Moreover, community standards cannot anticipate every possibility of novel research. There were no community standards about how to report neutrino observations which Davis could rely on to specify an appropriate level of externality. For Davis—and for the junior scientists in the schematic case—community standards will constrain their choices without fully determining what they should believe.

Here is a final attempt to reformulate the objection: One might hope that a philosophical analysis or theory of perception will determine the proper level of externality. Scientific observations could be posed at that level with only reference to epistemic values, and all other scientific claims would be inferences from claims at that base level.

If the default level of externality is not to be subject to revision or scrutiny in the course of inference, then it must favor security over significance to an extreme degree. After all, even a report of splodges presumes that there is actual equipment yielding readings and not facades or hallucinations. Insisting that scientists initially represent observations only in the most secure terms would yield scepticism or phenomenalist empiricism.⁵³

ch. 6, who considers numerous possibilities but concludes that there is no sensible way to construe scientific significance in purely objective or epistemic terms.

⁵³ Staley makes a similar point in relation to research teams. He writes that "if one were to ask that groups should ideally issue statements of group belief only when there is complete uniformity in what each individual member is 'compelled' to believe based on the evidence, one would in fact be saying that ideally such groups would issue almost no statements of any interest, and thus that there simply would be very little interesting empirical science" (Staley, "Evidential Collaborations," 328).

Perhaps the default level of externality need not be neutral sense data and may be theory-laden. All that the objection strictly requires is that scientists can rely on it regardless of utility considerations or value commitments. The objection still founders, because philosophical conceptions of evidence are little help in selecting a level of externality. To take one example, Peter Achinstein provides a theory of evidence which takes empirical data or phenomena as given.⁵⁴ It is about how bare observation becomes evidence, rather than about how observation becomes credible in the first place. I do not see how the problem is any less vexed on other theories of evidence.

To sum up: Selecting a level of externality requires balancing the desire for significant findings against aversion to mistakes. This balance of enthusiasm against caution depends on conditional utilities. That is, it is a matter of values. Attempts to see these values as somehow innocently epistemic fail.

A More Qualified Objection

In this section, I consider an argument by Bryce Huebner, Rebecca Kukla, and Eric Winsberg that values play *less* of a role in science like the search for the Higgs Boson than in policy-relevant research like climate modelling.⁵⁵ The result would, perhaps, be a limited version of the policy-relevance restriction.

Huebner et al. discuss the search for the Higgs boson and portray developments at CERN in a way that is initially congenial to my argument. They write that "inductive risk balancing continues to occur in unpredictable and, perhaps, unrecoverable ways throughout the research process, even where the research does not aim at some obviously value-laden goal."⁵⁶ Yet they go on to argue that the entanglement with values is importantly less complex for the discovery of the Higgs boson than it is for climate science. The existence or non-existence of the Higgs boson is a binary question. So, they argue, there are risks only along one dimension. Climate modelling does not involve a single binary question and so "researchers can have any of a wide, multi-dimensional array of

⁵⁴ Peter Achinstein, *The Book of Evidence* (Oxford: Oxford University Press, 2001).

⁵⁵ Bryce Huebner, Rebecca Kukla, and Eric Winsberg, "Making an Author in Radically Collaborative Research," in *Scientific Collaboration and Collective Knowledge*, eds. Thomas Boyer-Kassem, Conor Mayo-Wilson, and Michael Weisberg (Oxford: Oxford University Press, 2018): 95-116.

⁵⁶ Huebner et al., "Making an Author," 112.

investments in various outcomes...."⁵⁷ Arguing in this way, one might say that straight-forward existence questions are—although not value-free—less value-laden than more complex questions which matter to policy.

This difference is not as significant as Huebner et al. suggest. When an observation is construed at different levels of externality, the evidential context changes—that is, the observation is taken to tell us about different things. That means that what is presented as a binary question hides all sorts of other possibilities. Davis' question might have been posed as whether the dominant model of the sun was correct or not, but his observations were ultimately accepted as measurements of the rate of electron neutrinos arriving at Earth. So a claim can only be construed as a binary report against an implicit, multi-dimensional background of individual and community commitments.

Conclusion

The JRD thesis points to a tension at the heart of our epistemic lives. We aim to believe true things, and we aim to avoid believing false things. Our duties as knowers require us to arrive at some balance between these, and such a balance reflects value judgments: Conditional on the claim being true or false, what would the benefit or cost be of believing or not believing?

Neither the randomizer reply, the policy-relevance restriction, nor the lexical priority of evidence undo this. Values do not enter as merely a secondary consideration, nor do they enter only in situations with obvious policy consequences, nor do they enter only into inference.⁵⁸

⁵⁷ Huebner et al., "Making an Author," 112.

⁵⁸ An earlier version of this paper was presented at the Society for Exact Philosophy meeting in May 2018. Thanks to Chris Meacham and other interlocutors at the conference for helpful feedback.

EVIDENTIALISM, KNOWLEDGE, AND EVIDENCE POSSESSION

Timothy PERRINE

ABSTRACT: Evidentialism has shown itself to be an important research program in contemporary epistemology, with evidentialists giving theories of virtually every important topic in epistemology. Nevertheless, at the heart of evidentialism is a handful of concepts, namely evidence, evidence possession, and evidential fit. If evidentialists cannot give us a plausible account of these concepts, then their research program, with all its various theories, will be in serious trouble. In this paper, I argue that evidentialists has yet to give a plausible account of evidence possession and the prospects for doing so are dim.

KEYWORDS: evidentialism, evidence possession, knowledge, Richard Feldman, Earl Conee, Kevin McCain

In 1985, Earl Conee and Richard Feldman brought evidentialism into the limelight. At the core of their view was the following account of justification:¹

(EJ): An agent's doxastic attitude—belief, disbelief, suspended belief—towards a proposition p at a time t is justified if and only if having that doxastic attitude towards p fits the evidence the agent has (or possesses) at that time.²

Since then, evidentialism has been applied to many other issues, including the internalism/externalism debate,³ skepticism,⁴ epistemic value,⁵ epistemic

¹ This is an account of propositional justification. Conee and Feldman also develop a theory of “well-founded belief” which requires propositional justification and some additional conditions (see Earl Conee and Richard Feldman, “Evidentialism,” in *Evidentialism*, eds. Earl Conee and Richard Feldman (Oxford: Clarendon Press, 2004), 93-101.) What they call well-founded belief is closely related to what is sometimes called “doxastic justification.” In this paper, I will only be concerned with propositional justification.

² Cf. Conee and Feldman, “Evidentialism,” 83.

³ Earl Conee and Richard Feldman, “Internalism Defended,” in *Evidentialism*, eds. Conee and Feldman.

⁴ Richard Feldman, *Epistemology* (Upper Saddle River, New Jersey: Prentice Hall, 2003); Richard Feldman and Earl Conee, “Making Sense of Skepticism,” in *Evidentialism*, eds. Conee and Feldman.

⁵ Richard Feldman, “The Ethics of Belief,” in *Evidentialism*, eds. Conee and Feldman; Richard

Timothy Perrine

norms,⁶ defeaters,⁷ religious epistemology,⁸ and the epistemology of memory,⁹ among other things. Though evidentialism was initially offered as a theory of justified doxastic attitudes, it has become a research program.

At the center of this research program are a few key concepts: evidence, evidence possession, and evidential fit. Unfortunately, evidentialists have been less concerned with giving theories of these key concepts, perhaps because they think a positive feature of their theory is that it allows for different ways of spelling them out.¹⁰ But without accounts of these key concepts, evidentialism offers us not a *theory* but a *theory schema*.¹¹ Fortunately, evidentialists Richard Feldman and Kevin McCain have provided accounts of these concepts. Their accounts are interesting in their own right. But they are also of crucial importance for the success of evidentialism *qua* research program.

This paper critically evaluates those accounts, specifically their accounts of evidence possession. I begin, in section I, by reviewing Feldman's account of evidence and evidence possession. I argue that Feldman's account is much too restrictive to support the amount of knowledge humans possess. In section II, I review McCain's views of evidence and evidence possession. Like me, McCain finds Feldman's account too restrictive and aims for a moderate account. Nevertheless, in section III, I argue that McCain's account is open to several

Feldman, "Epistemological Duties," in *The Oxford Handbook of Epistemology*, ed. Paul Moser (Oxford: Oxford University Press, 2002).

⁶ Feldman, "The Ethics of Belief," Feldman, "Epistemological Duties," Trent Dougherty "The Ethics of Belief is (Just) Ethics," in *The Ethics of Belief*, eds. Jonathan Matheson and Rico Vitz (Oxford: Oxford University Press).

⁷ Richard Feldman, "Respecting the Evidence," *Philosophical Perspectives* 19, 1 (2005): 95-119; Trent Dougherty "Further Epistemological Considerations Concerning Skeptical Theism," *Faith and Philosophy* 28,3 (2011): 332-340.

⁸ Trent Dougherty, "Faith, Trust, and Testimony: An Evidentialist Reflection," in *Intellectual Virtue and Religious Faith*, eds. Timothy O'Connor and Laura Frances Goins (Oxford: Oxford University Press, 2014).

⁹ Earl Conee and Richard Feldman, "Evidence," in *Epistemology: New Essays* (Oxford: Oxford University Press, 2008), ed. Quentin Smith; Matthew Frise, "The Epistemology of Memory," *International Encyclopedia of Philosophy*.

¹⁰ Conee and Feldman suggest this at "Internalism Defended," 64 and "Evidence," 89.

¹¹ Or, for those sympathetic to the position, a "platitude" in desperate need of explication; cf. Trent Dougherty "Introduction" in *Evidentialism and Its Discontents*, ed. Trent Dougherty (Oxford: Oxford University Press, 2011).

counterexamples and some natural ways of amending his account have counterexamples as well.

But first a methodological remark. Evidentialism begun as an account of *justification*. But several philosophers—including William Alston,¹² Alvin Plantinga,¹³ and Richard Swinburne¹⁴—have worried that the term ‘justification’ does not pick out a single property, and thus there is no single property to give an account of. Though evidentialists do not necessarily fully embrace this conclusion, they do periodically defend their position by claiming that critics have misidentified the concept of justification at the heart of (EJ).¹⁵ Consequentially, counterexamples to (EJ) that turn on intuitions about whether a belief is “justified” or not are open to the criticism that the counterexamples turn on the wrong concept of justification. Fortunately, this is not the only way to give counterexamples to (EJ). For evidentialist usually insist that the kind of justification that (EJ) is about is the kind of justification that is necessary for knowledge; that is, evidentialists accept the following principle:

(KJ): A subject S knows that *p* at *t* only if S’s belief that *p* at *t* is justified.¹⁶

Consequently, in this paper, I will focus on whether evidentialists have provided a plausible account of a necessary condition for knowledge. By using

¹² William Alston, *Beyond “Justification”* (Ithaca: Cornell University Press, 2005).

¹³ Alvin Plantinga, “Justification in the 20th Century,” *Philosophy and Phenomenological Research* 50.1 (1990): 45–71.

¹⁴ Richard Swinburne, *Epistemic Justification* (Oxford: Oxford University Press, 2001).

¹⁵ Thus, Conee and Feldman (“Internalism Defended” 61–3) claim in response to Plantinga that they are not working with a deontic conception of justification; Conee and Feldman (“Postscript to ‘Evidentialism’” in *Evidentialism*, eds. Conee and Feldman, 103) claim in response to Fantl and McGrath that they do *not* use the phrase ‘justified in believing *p*’ as to imply that a subject has evidence sufficient for knowing *p*; Conee (“Postscript to ‘The Truth Connection’” in *Evidentialism*, eds. Conee and Feldman, 254–5) claims that there is a kind of epistemic justification that does not require evidence, but it is different from the kind of epistemic justification at the heart of (EJ); finally, Feldman (“Justification is Internal,” in *Contemporary Debates in Epistemology*, eds. Mattias Steup, John Turri, and Ernest Sosa (Oxford: Blackwell, 2013), 348) claims in response to Greco that the kind of epistemic justification he is interested in is not only different from justification understood as blameless believing but does not even require blameless believing.

¹⁶ Cf. Earl Conee “The Truth Connection,” in *Evidentialism*, eds. Conee and Feldman, 242; Conee and Feldman, “Internalism Defended,” 54; Conee and Feldman, “Evidence,” 83 fn. 1; Feldman, *Epistemology*; Kevin McCain *Evidentialism and Epistemic Justification* (London: Routledge, 2014); Dougherty, “The Ethics of Belief,” 159.

Timothy Perrine

examples that turn on cases of knowledge, not justification, we can bypass worries about identifying the wrong concept of justification.

I. Feldman on Evidence and Evidence Possession

According to Feldman, the evidence one possesses at a time is that subset of one's "total possible evidence" that meets the constraints of being both "available and acceptable." One's total possible evidence is all and only the information the person has "stored in his mind."¹⁷ This "storage" is meant to be quite inclusive, including both beliefs and experiences as well as both mental states one is currently thinking about and those one is not.¹⁸ Regarding the two constraints, Feldman spends almost no time on being "acceptable" except to say that the "acceptability" at issue is being *epistemically* acceptable and to criticize a simple account of it.¹⁹ Consequently, I'll set it aside. Regarding availability, Feldman argues that S has p available as evidence at t if and only if "S is currently thinking of p."²⁰ Since the evidence one possesses is the subset of total possible evidence available to one, for Feldman, the evidence one possesses at a time are those beliefs and non-belief states one is currently thinking about. Letting 'occurrent mental states' stand for the mental states (beliefs or otherwise) that one is thinking about, Feldman endorses:

Narrow View (NV): The evidence a subject S possesses at time t is the occurrent mental states S has at t.

(NV) is a highly restrictive theory of evidence possession. Consequently, there are many counterexamples to Feldman's view from cases of knowledge.²¹ After all, I know many things. For instance, I know that I'm a resident of China; that I am a brother; that I am more than 18 years of age; that the semester has just begun; that logical implication is transitive; etc. But perhaps just as obviously, I

¹⁷ Richard Feldman, "Having Evidence" in *Evidentialism*, eds. Conee and Feldman, 226.

¹⁸ Feldman, "Having Evidence," 232-41. According to Conee and Feldman, "Evidence," 87-88, experiences are "ultimate" evidence and beliefs "intermediate" evidence, but both are evidence.

¹⁹ Feldman, "Having Evidence," 226-7.

²⁰ Feldman, "Having Evidence," 232-41. Feldman ("Having Evidence," 240) suggest that for some propositions one can be currently thinking of them "non-consciously." I'm not sure that's possible, but as Feldman does not stress it, I don't take it to be an important suggestion.

²¹ The counterexamples here are similar to ones given by Alvin Goldman, "Internalism Exposed," *Journal of Philosophy* 96, 6 (1990): 271-93. But Goldman deploys them against an "accessibilism" position that is not logically equivalent to Feldman's.

know these things even if I'm not currently thinking about them, when for instance I'm in a dreamless sleep or I'm awake, but my attention is concerned with something other than those particular beliefs and evidence I may have for them. Thus, by (KJ), those beliefs are justified even while I sleep or my attention is otherwise preoccupied. But by (EJ) and (NV) it follows that those beliefs fit the occurrent mental states I have at those times. But clearly that, in general, will be false. Most of my beliefs, including the ones mentioned above, do not fit the occurrent mental states I have at any given time. Thus, most of the time most of my beliefs will not constitute knowledge. But that is an absurd result. Feldman's view, while falling short of skepticism about knowledge and justification, comes too close to it.

In response to these kinds of cases, Feldman might claim that there are occurrent and dispositional senses of 'knows' and cognates.²² Thus, when not considering the evidence I have for the proposition (e.g.) that I am a Chinese resident I might still be said to "dispositionally know" that. Even without fussing over what exact account of "dispositional knowledge" to give, we can see that there are two problems with this proposal.

First, this response rests on their being a distinction between a dispositional and occurrent sense of 'knows.' But there is no independent reason for thinking there is such a distinction. Indeed, there is reason for doubting that there is such a distinction. For our ordinary practice of attributing knowledge is usually insensitive to facts about the experiences of others at the time of attribution. Learning what experiences a person was (or wasn't) undergoing at t , after having attributed knowledge that p to her at t , does not usually result in a change or modification of our attribution of knowledge.²³ That's certainly not what we should expect if there were such a distinction.

²² Compare Feldman, "Having Evidence," 237; Conee and Feldman, "Internalism Defended," 67-8.

²³ The only exception I can think of concerns the acquisition of knowledge. We sometimes attribute knowledge to a person at a time because we believe the person had an experience at that time that is responsible for them acquiring the knowledge we attribute to them. If we learned they did not have that experience, we would retract our attribution of knowledge. But clearly this exception does not help Feldman. After all, even in this kind of case, we do not go from attributing one kind of knowledge to another, but from attributing knowledge to ignorance. Worse yet, the problem cases for Feldman do not concern acquiring knowledge, but knowledge already possessed.

Second, even granting the distinction, this response faces a dilemma. Either dispositionally knowing something implies knowing it or it does not. If it does not imply knowing, then this is not a response at all, it merely gives a label to the problem. For, on it, it still comes out that I usually do not know that (e.g.) I am a Chinese resident. Suppose, by contrast, dispositionally knowing something implies knowing it. On this response, then, knowledge is bifurcated: S knows that p if and only if S either occurrently knows that p or dispositionally knows that p . But this horn of the dilemma requires a rejection of (KJ). After all, since most of the time most of my beliefs do not fit the evidence I possess that is given by my occurrent experiences, it follows by (EJ) that most of the time most of my beliefs are unjustified. But, nonetheless, many of those unjustified beliefs constitute knowledge, namely, dispositional knowledge. Since this horn requires the rejection of (KJ) it is safe to assume most evidentialists would not prefer it.

In the face of these difficulties, Feldman seems most inclined to bite the bullet, and embrace a kind of (moderate) skepticism.²⁴ But there's no reason to bite the bullet here; it is much more likely that we've simply taken a wrong turn somewhere, presumably at Feldman's overly restrictive account of evidence possession.

II. McCain's Moderate View

Like me, McCain finds Feldman's theory of evidence possession to be overly restrictive and implausible.²⁵ McCain aims to provide a more moderate position. McCain distinguishes two camps on the ontology of evidence. According to *Psychologism*, evidence only consists of psychological items, specifically, one's non-factive mental states.²⁶ (Non-factive mental states are representational mental states that "one can be in even if they misrepresent the word."²⁷) According to *Anti-Psychologism*, evidence only consists of non-psychological items.²⁸ Among the latter camp, McCain draws a further distinction. According to

²⁴ Compare Feldman "Having Evidence," 237.

²⁵ Conee also seems sympathetic to a more moderate view, but does not develop one in the detailed way McCain does.

²⁶ McCain, *Evidentialism and Epistemic Justification*, 10 fn. 5.

²⁷ McCain, *Evidentialism and Epistemic Justification*, 10.

²⁸ McCain, *Evidentialism and Epistemic Justification*, 10. Notice that, so defined, while *Psychologism* and *Anti-Psychologism* are mutually exclusively, they are not mutually exhaustive.

Propositionalism, evidence only consists of propositions. (McCain assumes that propositions are non-psychological items.²⁹) Finally, a sub camp of *Propositionalism* is *Factive-p*: evidence only consists of true propositions.³⁰

McCain argues against *Factive-p*. Beyond that, he is neutral with regard to *Psychologism* and *Propositionalism*. He recognizes that the two positions offer incompatible ontologies for evidence but “the disagreement does not lead to significant epistemic differences.”³¹ McCain’s point seems plausible when it comes to the issue of evidence possession. For instance, on *Psychologism*, the evidence I possess might be my mental state of believing that *p*; but on *Propositionalism*, the evidence is the proposition *p*, and the reason why I possess it is because I stand in the “believing” relation to it. Such a difference does not seem so grand. However, for easy of exposition, in what follows I’ll frequently write as if it is mental states that provide evidence and not the propositional content of those mental states.

Regarding evidence possession, whereas Feldman thought of it as a two-place relation between a person and a body of evidence, McCain suggest that it should be thought of as a three place relation between a person, a body of evidence, and a proposition.³² As he sees it, one doesn’t just “have” evidence; one has evidence for/against a proposition. Regarding the accessibility of evidence, McCain considers the position that the evidence one possesses is one’s total possible evidence, i.e. *any* information stored in one’s mind. Let’s call that view:

Wide View (WV): The evidence a subject *S* possesses at time *t* is any and all information stored by *S* at *t*.

McCain, following Feldman, rejects (WV) as too permissive. He gives the following counterexample:

DEEP MEMORY: Sara is a normal adult in her thirties. Sara has many memories of her childhood that she can recall. Some of these memories she can easily recall

²⁹ McCain, *Evidentialism and Epistemic Justification*, 21.

³⁰ McCain, *Evidentialism and Epistemic Justification*, 10-11. McCain identifies another position he calls *Non-Factive_p* according to which “evidence consists only of propositions, but those propositions can be true or false” (*Evidentialism and Epistemic Justification*, 11). However, *Propositionalism* and *Non-Factive_p* would seem to differ only if there could be propositions which were not true or false, in which case *Propositionalism* would allow them to be evidence, and *Non-Factive_p* would not. As McCain never really discusses this possibility, I take it that there’s not really an important difference between *Propositionalism* and *Non-Factive_p*.

³¹ McCain, *Evidentialism and Epistemic Justification*, 27.

³² McCain, *Evidentialism and Epistemic Justification*, 49-50.

Timothy Perrine

and some she can only recall with prompting of specific kinds. One particular memory, that it was raining on the third day of March when Sara was three years old, is very deeply stored. Sara could only bring this memory to consciousness with years of training and psychological therapy. At t Sara has not undergone any of the training or psychological therapy.³³

McCain thinks it is implausible that Sara's memory is evidence that she possesses. For if it were, then she would be justified in believing that it was raining on that day; but intuitively she is not. Consequently, the evidence one has cannot be one's *total possible evidence*; (WV) is false.

McCain aims for a moderate account that is more inclusive than Feldman's narrow account but not as permissible as the wide account. To that end, he proposes:

Moderate View (MV): S has p available as evidence relevant to q at t iff at t S is currently aware of p or S is disposed to bring p to mind when reflecting on the question of q 's truth.³⁴

Given McCain's neutral attitude between *Psychologism* and *Propositionalism*, his p can range over either non-factive mental states (per *Psychologism*) or the propositional content of those non-factive mental states (per *Propositionalism*) though not both. Since my criticisms of McCain focus mainly on evidence possession, I'll ignore this complication.

Now strictly speaking (MV) is an account of *available* evidence and not evidence *possession*. But McCain, following Feldman, holds that the evidence one possess is that subset of one's total evidence that is both available and "epistemically acceptable."³⁵ However, McCain rarely touches on this second condition, and his informal gloss on it seems to amount to little more than that there be no counterexamples to the theory of evidence possession. Consequently, in describing McCain's view, I'll follow his lead and freely move between talk of available evidence and evidence possession.

III. Problems with McCain's Account

McCain's view of evidence possession can handle some counterexamples to Feldman's account. For it allows beliefs I am not currently thinking of to be part of

³³ McCain, *Evidentialism and Epistemic Justification*, 35; cf. Feldman, "Having Evidence," 228-9.

³⁴ McCain, *Evidentialism and Epistemic Justification*, 51.

³⁵ McCain, *Evidentialism and Epistemic Justification*, 34.

the evidence I possess; consequently those beliefs can strongly support other beliefs like (e.g.) I am a Chinese resident or I'm not 18 years old. Nevertheless, McCain's view is still inadequate. I'll argue that there are counterexamples both to it and natural ways of revising it.

First, there are counterexamples to McCain's view because it ties evidence possession too closely to one's dispositions and what a person is disposed to think of when that person considers the truth of a proposition need not be the evidence that person possesses. Consider:

RAY. Ray is a racist, who is nevertheless a leading scholar on cognitive development in children. Ray has always thought that members of a certain race were none too bright. Early in his career, Ray has performed and published numerous studies which conclusively support the conclusion that children of a certain race develop more slowly than others on certain skills. Ray remembers those studies, and can summarize his findings if requested. Nevertheless, Ray himself rarely thinks about his studies. Further, when he reflects as to why children of a certain race develop more slowly than others, he is almost never disposed to consider his studies but rather his racist reasons for the belief.

Because Ray is not disposed to bring to mind his studies when he thinks about their conclusions, by (MV), his memories of his studies are not part of his evidence for their conclusions. But that is very unintuitive. Ray, after all, has done numerous studies and could, if asked, summarize them. Further, it would be entirely appropriate to use Ray as (say) an expert witness at a trial not only because of his status as a leading scholar but also because, it seems, he has such excellent evidence for the conclusions of his studies that he could provide for a jury. Of course, it may be that Ray is not justified in believing his conclusions; perhaps the fact that his beliefs are casually sustained by poor reasons is sufficient for his beliefs in the conclusion of his studies to be unjustified.³⁶ But that is consistent with my point that Ray's memories of his studies are part of his evidence for those conclusions; they should not be demoted out of his possessed evidence for those conclusions just because he isn't disposed to bring them to mind when considering those conclusions.

While Ray is able to bring to mind his studies, he isn't disposed to. This suggests we offer a weaker account than McCain's official one as follows:

³⁶ Or, for those who like the distinction between "propositional" and "doxastic" justification/well-founded belief, perhaps Ray's beliefs are propositionally justified, but not doxastically justified/well-founded.

(MV*): S has p available as evidence relevant to q at t iff at t S is currently aware of p or S has the ability to reflect on the question of q 's truth and bring p to mind when so reflecting.

(MV*) avoids the problem of RAY. Further, it remains a *moderate* view, as it includes among one's evidence more than one's occurrent mental states but does not include Sara's inaccessible memory in DEEP MEMORY because she is unable to bring it to mind. Indeed, (MV*) is even suggested by some of McCain's informal remarks, such as "...stored information is available as evidence on a particular topic when S can recall this information by reflecting on the topic."³⁷

The problem with (MV*) is that it over-intellectualizes evidence possession because some cognizers—including animals, small children, and mentally handicapped adults—have non-occurrent beliefs that constitute knowledge but lack the ability to willfully reflect on the truth of their beliefs. Consider:

HAL: Due to various mental handicaps, Hal's attention span is extremely small and he is unable to reflect or otherwise follow a line or train of thought. Hal believes that his sister Monique lives in Nashville. He's visited her house many times and has many memories of his visits. If asked where she lives, he will consistently (i) bring to mind those memories and (ii) respond that she lives in Nashville. At time t , though, Hal is thinking about something else.

At t , Hal is unable to reflect on the truth of his belief. By (MV*), the only non-occurrent mental states that are part of Hal's evidence for his belief are those that he'd bring to mind upon reflecting. Thus, by (MV*), it follows that none of his non-occurrent mental states—including his memories—are part of his evidence for his belief about where his sister lives. That is an implausible result. But things are worse. Given (MV*), the only evidence Hal has for where his sister lives is his occurrent mental states. But since they do not concern his belief about where his sister lives, his belief that his sister lives in Nashville does not fit the evidence he has at t . By (EJ), it follows that the belief is not justified at t , and by (KJ) that he does not know it. But intuitively Hal does know where his sister lives at t . For these reasons (MV*) should be rejected.

That counterexample shows that the possession of evidence should not be tied too closely to a cognizer's ability to reflect, on their own volition, about the truth of a belief. Thus, we might try weakening (MV*) to get:

(MV**): S has p available as evidence relevant to q at t iff at t S is currently aware

³⁷ McCain, *Evidentialism and Epistemic Justification*, 50.

of p or if S 's attention were directed to q , then S would be able to bring p to mind.

(MV**) is immune from cases like HAL. For even if Hal is unable to, on his own, direct his attention to a proposition, it does not follow that his attention could not be so directed. And all that the second disjunct in (MV**) requires is that *were* Hal's attention directed to the relevant proposition (e.g. "my sister lives in Nashville") *then* Hal has the ability to bring his memories to mind, which of course he does. (MV**) also handles RAY. For if Ray's attention were directed to the relevant proposition, he is able to bring to mind his studies, even if he's not disposed to. Finally, (MV**) is a moderate position because it excludes Sara's specific memory in DEEP MEMORY as being part of her possessed evidence because she lacks the ability to bring it to mind.

However, there are counterexamples to (MV**). These counterexamples involve what I'll call *evidentially isolated basic beliefs*. A basic belief is, roughly, a belief that constitutes knowledge independent of its positive epistemic relations to other beliefs.³⁸ S 's belief that p at t is an evidentially isolated belief just when there are no other non-factive mental state (or states) m such that (A) S has m at t , and (B) m supports p such that given just m believing p is the doxastic attitude that "fits." Now from the mere fact that a belief is a basic belief it does not necessarily follow that it is also evidentially isolated. Even if (e.g.) I can know that p in a basic way because you testified that p , it may also be the case that I have other beliefs that support p or beliefs that support that *if* you were to testify that p , then p is very likely to be true.³⁹ However, an evidentially isolated basic belief would be a belief that is both a basic belief and also evidentially isolated. More formally: S 's belief that b is an evidentially isolated basic belief at time t if and only if (i) b is a basic belief at t and (ii) there are no other non-factive mental state (or states) m such that (A) S has m at t , (B) m supports b such that given just m believing b is the doxastic attitude that "fits."

³⁸ Sometimes (e.g. Alvin Plantinga, *Warrant the Current Debate* (New York: Oxford University Press, 1993)) these are called "properly basic beliefs." There are several well-known defense of the existence of basic beliefs in the literature.

³⁹ Of course, not everyone agrees that testimony is a basic source of knowledge. See Jennifer Lackey, *Learning from Words*, (Oxford: Oxford University Press, 2008) or Elizabeth Fricker, "Against Gullibility," in *Knowing From Words*, eds. Bimila Krishna Matilal and Arindam Chakrabarti (Dordrecht: Springer, 1994) for criticism of that position. I criticize Lackey's argument in Timothy Perrine, "In Defense of Non-Reductionism in the Epistemology of Testimony," *Synthese* 191, 14 (2014): 3227-3237.

Evidentially isolated basic beliefs, so defined, would provide counterexamples to (MV**). Suppose at time t S's belief that b is an isolated basic belief. By definition, at t S's belief that b constitutes knowledge. From (KJ), it follows that at t S's belief that b is justified. From (EJ), it follows that at t S's belief that b fits the evidence that S possesses at t . And, from (EJ) and (MV**), it follows that (1) there is some mental state (or states) m such that S is currently aware of m or if S's attention were directed to b , then S would bring m to mind, and (2) given m believing b is that doxastic attitude that "fits". But, by definition of b being an isolated basic belief, (1) and (2) do not hold. For, by definition, S does not have any mental states s such that given s believing b is the doxastic attitude that fits. Thus, S does not have any mental state s such that S is current aware of s or if S attention were direct to b then S would bring s to mind which is also such that given s believing b is the doxastic attitude that fits. Thus, given the existence of evidentially isolated basic beliefs, as well as the principles (KJ) and (EJ), there are counterexamples to (MV**). The interesting question is thus whether there are any evidentially isolated basic beliefs.

It is plausible that there are. Consider the following two cases.

BIRD WATCHER. While hiking in a mountain range, Emmett an expert bird watcher sees what might be a rare bird up ahead. After positioning himself with a clear view of the bird, he immediately identifies it as a male goldfinch—a bird he has seen many times, but is not known to be in this mountain range. Upon the basis of his visual experience, he immediately forms the belief that there is a male goldfinch in the woods. Pleased with his observation, and tired from his hike on the mountain range, Emmett returns to his camp where he takes a nap.

LOGIC. While reading ahead in her logic textbook, Sidra considers for the first time whether the conjunction elimination rule in her logic textbook is sound. It seems overwhelming obvious to her that it must be sound, and she comes to believe that it is. After completing her homework, she plays a serious game of volleyball with some of her friends.⁴⁰

⁴⁰ This case is modelled on one given in Andrew Moon, "Knowledge without Evidence," *Mind* 121, 482 (2012): 309-331. Moon criticizes a logically distinct and stronger position than me: that S's knowledge that p requires S believes that p on the basis of evidence, E, and further, S can be aware of that evidence E by way of introspection at t . My criticisms of evidentialism have not relied upon claims about based evidence or introspection. Further, Moon does not bring out what I take to be most important in these examples: that the basic beliefs are *evidentially isolated*; in fact, if we add to his case that the belief is not evidentially isolated, his counterexample would fail. Thus, I take my discussion to extend, if not supplant, his.

While he sleeps, Emmett's belief that he saw a male goldfinch in the woods is an evidentially isolated basic belief. Presumably, it not only constitutes knowledge while Emmett sleeps but is known in a basic way. Does Emmett have any other mental states while he sleeps that could be possessed evidence for his belief that he saw a male goldfinch in the woods? Clearly, whatever occurrent mental states Emmett has while he sleeps do not support that belief. Thus, given (MV**), if Emmett has any evidence for his belief while he sleeps it must be other mental states that he has (non-occurrently) while he sleeps that he would bring to mind, were his mind directed to the proposition that he saw a male goldfinch in the woods that day. But it is doubtful that Emmett has such other mental states.⁴¹ After all, Emmett does not antecedently believe that there are goldfinches in this forest; after all, he knows that goldfinches do not generally inhabit this mountain range. And it is hard to see that there are other beliefs Emmett formed when he formed the belief that he saw a male goldfinch in the woods that would support that belief to such a degree that they would make such a belief justified on their own. Thus, given (MV**), while he sleeps, Emmett does not have evidence for his belief and, thus, from (EJ) and (KJ) he does not know that he saw a male goldfinch in the woods, despite what is most intuitive.

Similar points apply to Sidra's belief that conjunction elimination is sound while she plays volleyball with her friends. That belief is presumably a basic belief. It is also an evidentially isolated basic belief. Though Sidra is having occurrent experiences as she plays volleyball, clearly none of them are evidence for the belief that conjunction elimination is sound. Thus, given (MV**), if she has any evidence for her belief while she plays volleyball, it must be other mental states that she has (non-occurrently) while she is playing that she would bring to mind, were her mind directed to the proposition that conjunction elimination is sound. But it is doubtful that she has other such beliefs. (Sidra is, after all, a student not a logic professor.) Thus, given (MV**), while she plays volleyball, Sidra does not have any evidence for her belief and thus, from (EJ) and (KJ), she does not know that conjunction elimination is sound, despite what is most intuitive.

In response, McCain might claim that Emmett does have evidence: namely a "disposition to recollect" that there is a male goldfinch in the forest, where a disposition to recollect something is "a disposition to bring to mind the proposition as known."⁴² But this response is unsuccessful. Here is a dilemma argument against

⁴¹ Or that we must understand the example in this way.

⁴² McCain, "No Knowledge without Evidence," *Journal of Philosophical Studies* 40 (2015): 369–

it. Either a disposition to bring to mind the proposition as known is (i) a non-occurrent belief that Emmett has while he sleeps that states that he knows that he saw a goldfinch in woods that day or (ii) it is a disposition to form such a belief that has yet to manifest. If (i), then that non-occurrent belief may very well be evidence that Emmett has for his belief. But there's no reason to think that Emmett has formed that belief or that he must. For such a belief is actually a higher-order belief—a belief about another belief—and there's no reason to suppose that when we form simple perceptual beliefs we also thereby form higher-order beliefs about those simple perceptual beliefs. If (ii), then it may be plausible that Emmett has the relevant disposition. However, the relevant disposition is not evidence because it is not a representation of the world,⁴³ but (at best) a disposition to represent the world. But if it is not a representation, then it can't be a non-factive mental state or the propositional content of a non-factive mental state, since those are representations. But given that McCain is committed to *Psychologism* or *Propositionalism*, it follows that a disposition to recollect, so understood, couldn't be evidence. So either Emmett lacks the mental state that could be evidence or has a disposition that couldn't be evidence. Either way, Emmett does not have evidence while he sleeps.

Here is related objection. It may be that cognizers like Emmett and Sidra have various dispositions such that were those disposition to manifest they would provide mental states that could serve as evidence for Emmett's and Sidra's beliefs. For instance, on some views, if it seems to you that *p*, then you thereby have evidence that *p*. Further, if Sidra and Emmett were to reflect on their beliefs, perhaps it would seem to them that, respectively, Emmett saw a goldfinch in the woods that day and conjunction elimination is a sound rule. However, these points even if true do not undermine my argument. For even if Emmett and Sidra could acquire evidence by reflecting and manifesting various dispositions, they do not yet have that evidence for they have yet to do the relevant reflecting. Thus, they do not yet possess the evidence that they might acquire through such reflecting. Further, the disposition to form something with propositional content that could

76. McCain is quoting Earl Conee and Richard Feldman, "Replies," in *Evidentialism and Its Discontents*, ed. Dougherty, 304. McCain gives this response to Moon, "Knowledge without Evidence."

⁴³ McCain (*Evidentialism and Epistemic Justification*, 11) lists what he considers the relevant kinds of non-factive mental states, but they constitute beliefs, experiences and "perhaps others such as intuitions and rational insights." He does not include dispositions nor should he.

be evidence does not itself have propositional content. (Compare: a disposition to yell is not air vibrations but, when manifested, does produce air vibrations.) So the dispositions themselves do not count as evidence that Emmett has while he naps or Sidra possesses while she plays volleyball.

We can generalize the points of the previous paragraphs as follows. I've already shown that given the existence of evidentially isolated beliefs, as well as (KJ) and (EJ), it follows that (MV**) is false. Thus, for these responses to undermine my argument they must show that these kinds of cases could not be understood as cases of evidentially isolated basic beliefs. The most promising way to do that is to show that Emmett's and Sidra's beliefs are not evidentially isolated. To show that their beliefs are not evidentially isolated, one must identify some non-factive mental state (or states) that (A) they have at the relevant time but also (B) support the relevant belief so that given just that mental state belief is the relevant doxastic attitude that fits. The problem with these attempts is that they fail to find a mental state that satisfies both (A) and (B). Some mental states—like non-occurrent beliefs or occurrent experiences like seemings—may satisfy (B), but there is no reason for assuming that (A) must also always be met in these cases for those mental states. Some dispositions to form mental states meet (A) in these cases but there is no reason for thinking that (B) is met with regard to them.

To be sure, this criticism of (MV**) requires not just the existence of basic beliefs but the existence of evidentially isolated basic beliefs. But that does not strike me as overly objectionable. For it is plausible that people do have evidentially isolated basic beliefs. I've given two plausible examples of such beliefs—one concerning a past event, one concerning a simple logical truth. But even if one is not convinced by those particular examples, it is plausible that at least some of the things we know about past events and simple logical truths are evidentially isolated basic beliefs.⁴⁴

⁴⁴ Might the evidentialist eschew basic beliefs altogether, urging instead a kind of coherentism about justification? In response, even if an appeal to coherentism would deliver sufficient justification for knowledge about simple logical truths and past events, my criticism would still show something noteworthy: that evidentialism must be developed as to take a side on the foundationalism/coherentism/infinitism dispute. But I doubt that coherentism will save the evidentialist from the problem of evidentially isolated beliefs. For coherentists usually require that a belief cohere with a set of beliefs, where "cohere" means more than logical consistence but includes things like probabilistic consistent and explanatory relevance (cf. Laurence Bonjour, *The Structure of Empirical Knowledge* (Cambridge: Harvard University Press, 1985)). But if a belief really is evidentially isolated, then it is doubtful it will "cohere" in this sense with other

The problem of evidentially isolated basic beliefs is different from the objection of forgotten evidence that is periodically pressed against evidentialism.⁴⁵ Consider two cognizers, Sally and Sid. Both Sally and Sid form the same belief, say, broccoli has health benefits. But Sally forms this belief on the basis of another belief that is good evidence for it, say, that a New York Times science article reports as much. By contrast, Sid forms this belief on the basis of another belief that is not good evidence for it, say, that a National Enquirer article reports as much. Finally, suppose at a later time, both forget their evidence, i.e. their beliefs about the New York Times and National Enquirer, but do not acquire any new evidence for these beliefs. Intuitively, this objection goes, Sally is justified in her belief, even though she has forgotten her evidence. By contrast, Sid is not justified in his belief, even if he mistakenly believes that he did form the belief in an epistemically appropriate way. But, then, at this later time, it is not just the evidence that Sally and Sid possess at that time that is relevant to the justification of their beliefs, as (EJ) would have. Rather, the evidential strength of the beliefs they initially had and used to acquire the belief are also relevant, even though they have since forgotten that evidence; as Goldman put it once, “earlier evidence is also relevant to justifiedness.”⁴⁶ So, the objection goes, (EJ) is false.

My objection is distinct from that objection. Specifically, that objection claims that the evidential strength of the beliefs one uses to acquire a belief are relevant to the justificatory status of the acquired belief, even if one forgets one’s initial evidence. But my criticism does not turn on this claim. Indeed, I can concede that for any non-basic belief, that belief is justified at a time if and only if that belief fits the evidential strength of other beliefs that person has at that time. If a person formed a non-basic belief on the basis of good evidence, but at a later time forgot that good evidence and acquired no new beliefs that were equally good or better evidence as the old beliefs, then at the subsequent time the person’s belief is not justified. Those concessions are consistent with what is needed for my argument: that there are evidentially isolated basic beliefs.⁴⁷

beliefs and thus could be justified by cohering with other beliefs.

⁴⁵ See Alvin Goldman, “Internalism Exposed,” *Journal of Philosophy* 96, 6 (1999): 271-93 for an influential presentation, which I follow. For a similar case, see John Greco, “Internalism and Epistemically Responsible Belief,” *Synthese* 85, 2 (1990): 245-77.

⁴⁶ Alvin Goldman, “Toward A Synthesis of Reliabilism and Evidentialism?” in *Evidentialism and Its Discontents*, ed. Dougherty, 267.

⁴⁷ To be sure, one could *develop* the forgotten evidence objection into an objection similar to mine. For instance, one might argue that in LOGIC Sidra’s experience as of conjunction

Finally, notice that a retreat to either the narrow view of evidence possession (NV) or the wide view (WV) will not help this problem either. Since (NV) is more restrictive than moderate views like (MV)-(MV**), it cannot help. But (WV) cannot either, despite being a more permissive account. For an evidentially isolated belief is one that, by definition, is isolated from the *rest* of one's beliefs, including the ones that a person cannot access. Thus the problem of evidentially isolated basic beliefs is a serious problem for any of these ways that an evidentialist might develop his account.

IV. Conclusion

Evidentialism is an important research program in contemporary epistemology. At the heart of that research program are a few key concepts: evidence, evidential fit, evidence possession. In this paper, I've argued on the basis of various examples that evidentialists have yet to provide us with a plausible theory of evidence possession. Consequently, the success of their research program is drawn into question until they do so.⁴⁸

elimination being sound was itself very good evidence for believing conjunction elimination is sound and that her belief is a basic one because it was formed on the basis of such evidence. But notice (i) as a matter of fact, neither Goldman nor others do develop the objection this way, and (ii) it would be misleading to speak of this as an objection from forgotten evidence since experiences are not forgotten, beliefs are.

⁴⁸ For helpful feedback and comments, I thank Jordi Cat, Dave Fisher, Hao Hong, Mark Kaplan, Tufan Kiyamaz, Tim Leisz, Adam Leite, Kevin McCain, Nick Montgomery, Timothy O'Connor, Luis Oliveira, Harrison Waldo as well as an audience at Indiana University and the 2015 Indiana Philosophical Association.

HIGHER-ORDER DEFEAT WITHOUT EPISTEMIC DILEMMAS

Mattias SKIPPER

ABSTRACT: Many epistemologists have endorsed a version of the view that rational belief is sensitive to higher-order defeat. That is to say, even a fully rational belief state can be defeated by (sufficiently strong) misleading higher-order evidence, which indicates that the belief state is irrational. In a recent paper, however, Maria Lasonen-Aarnio calls this view into doubt. Her argument proceeds in two stages. First, she argues that higher-order defeat calls for a *two-tiered* theory of epistemic rationality. Secondly, she argues that there seems to be no satisfactory way of avoiding *epistemic dilemmas* within a two-tiered framework. Hence, she concludes that the prospects look dim for making sense of higher-order defeat within a broader theoretical picture of epistemic rationality. Here I aim to resist both parts of Lasonen-Aarnio's challenge. First, I outline a way of accommodating higher-order defeat within a *single-tiered* framework, by amending epistemic rules with appropriate provisos for different kinds of higher-order defeat. Secondly, I argue that those who nevertheless prefer to accommodate higher-order defeat within a two-tiered framework can do so without admitting to the possibility of epistemic dilemmas, since epistemic rules are not always accompanied by 'oughts' in a two-tiered framework. The considerations put forth thus indirectly vindicate the view that rational belief is sensitive to higher-order defeat.

KEYWORDS: higher-order defeat, higher-order evidence, epistemic dilemmas, epistemic rules, epistemic rationality

1. Introduction

Many epistemologists have endorsed a version of the view that rational belief is sensitive to higher-order defeat. That is to say, even a fully rational belief state can be defeated by (sufficiently strong) misleading higher-order evidence, which indicates that the belief state is irrational. Here is a putative example:¹

¹ Similar cases of misleading higher-order evidence can be found in, e.g., David Christensen, "Higher-Order Evidence," *Philosophy and Phenomenological Research* 8 (2010), 185-215, and "Formulating Independence," In *Higher-Order Evidence: New Essays*, eds. Mattias Skipper and Asbjørn Steglich-Petersen (Oxford: Oxford University Press, forthcoming); Kevin Dorst, "Higher-Order Uncertainty," in *Higher-Order Evidence: New Essays*, and "Evidence: A Guide for the Uncertain," *Philosophy and Phenomenological Research* (forthcoming); Sophie

Self-Enhancement Bias: John rationally believes that he is better than most people at driving. However, when reading today's newspaper, John learns about the well-documented *self-enhancement bias*: the widespread tendency to overrate oneself on a wide range of qualities and abilities, including intelligence, driving skills, and so on. As it happens, John is one of the few people who does not suffer from the bias.

Upon learning about the self-enhancement bias, how, if at all, should John revise his opinion about his own driving skills? In order to avoid getting sidetracked by issues concerning the (un)reliability of current journalism and experimental psychology, let us simply assume that John, after having read the newspaper, has good reason to think that the self-enhancement bias is indeed a widespread and pervasive phenomenon. Given this, it seems plausible to say that John should give up his belief (or at least become less confident) that he is better than most people at driving. After all, he has strong—albeit misleading—reason to think that his belief is the result of an irrational bias, and it seems plausible to say that one should give up beliefs that one has strong reason to think are irrational.

In recent years, a number of attempts have been made at vindicating the intuition that agents like John should revise their opinions, even if the higher-order evidence at hand is misleading. Perhaps the most prominent suggestion is due to David Christensen,² who points out that someone who disregards a body of misleading higher-order evidence seems to fall prey to an objectionable kind of dogmatism or question-begging reasoning: if John does not revise his opinion in response to the evidence about the self-enhancement bias, he must, it seems, take the evidence to be misleading. But the evidence is obviously only misleading if John does not suffer from the self-enhancement bias. So if John does not revise his opinion in response to the evidence about the self-enhancement bias, he must treat himself as an exception to the rule. However, the thought goes, John thereby

Horowitz, "Epistemic Akrasia," *Noûs* 48 (2014): 718–44; Mattias Skipper, "Reconciling Enkrasia and Higher-Order Defeat," *Erkenntnis* (forthcoming), and "Higher-Order Defeat and the Impossibility of Self-Misleading Evidence," in *Higher-Order Evidence: New Essays*; Daniel Whiting, "Against Second-Order Reasons," *Noûs* 51 (2017): 398–420, and "Whither Higher-Order Evidence?," in *Higher-Order Evidence: New Essays*; Alex Worsnip, "The Conflict of Evidence and Coherence," *Philosophy and Phenomenological Research* 96 (2018): 3–44, and "How Your Total Evidence Can Mislead About Itself," in *Higher-Order Evidence: New Essays*.

² David Christensen, "Epistemology of Disagreement: The Good News," *Philosophical Review* 116 (2007): 187–217; "Higher-Order Evidence"; "Disagreement, Question-Begging, and Epistemic Self-Criticism," *Philosophers' Imprint* 11 (2011); "Formulating Independence."

seems to beg the question the against the studies on the self-enhancement bias. Contraposing: since John *shouldn't* engage in this kind of question-begging reasoning, he *should* revise his opinion about his own driving skills in reponse to the evidence about the self-enhancement bias.³

Despite the intuitive pull of this diagnosis, Maria Lasonen-Aarnio⁴ has recently challenged the view that rational belief is sensitive to higher-order defeat. Her argument proceeds in two stages. First, she argues that higher-order defeat calls for a *two-tiered* theory of epistemic rationality. Secondly, she argues that there seems to be no satisfactory way of avoiding *epistemic dilemmas* within a two-tiered framework. Hence, she concludes that the prospects look dim for making sense of higher-order defeat within a broader theoretical picture of epistemic rationality. If she is right, we are faced with the counterintuitive result that John should *not* revise his opinion about his own driving skills upon learning about the self-enhancement bias (and, more generally, that that a rational belief can remain rational even in light of strong evidence to the contrary).

Here I aim to resist both parts of Lasonen-Aarnio's challenge. After laying out Lasonen-Aarnio's argument in more detail (§2), I outline a way of accommodating higher-order defeat within a *single-tiered* framework, by amending epistemic rules with appropriate provisos for different types of higher-order defeat (§3). I then argue that those who nevertheless prefer to accommodate higher-order defeat within a two-tiered framework can do so without admitting to

³ For other considerations in favor of the view that rational belief is sensitive to higher-order defeat, see also Miriam Schoenfield, "An Accuracy Based Approach to Higher Order Evidence," *Philosophy and Phenomenological Research* 93 (2016): 1-26; Adam Elga, "Reflection and Disagreement," *Noûs* 41 (2007): 478-502. In effect, Elga argues that *peer disagreement* must act as a higher-order defeater, since one could otherwise become justified in considering an epistemic *peer* to be an epistemic *inferior* merely by discovering that he or she disagrees with oneself. While many have found Elga's view basically correct, those who don't typically still accept that peer disagreement can have defeating force—see, e.g., Thomas Kelly, "Peer Disagreement and Higher Order Evidence," in *Social Epistemology: Essential Readings*, eds. Alvin I. Goldman & Dennis Whitcomb (New York: Oxford University Press, 2010), 183-217; Jennifer Lackey, "What Should We Do When We Disagree?," In *Oxford Studies in Epistemology*, 3, eds. Tamar Szabó Gendler and John Hawthorne (Oxford: Oxford University Press, 2008), 274-93. To my knowledge, Michael G. Titelbaum, "Rationality's Fixed Point (Or: In Defense of Right Reason)," In *Oxford Studies in Epistemology* 5 (2015) is the only current proponent of the view that disagreement can never have defeating force.

⁴ Maria Lasonen-Aarnio, "Higher-Order Evidence and the Limits of Defeat," *Philosophy and Phenomenological Research* 88 (2014): 314-45.

the possibility of epistemic dilemmas, since epistemic rules are not always accompanied by 'oughts' in a two-tiered framework (§4). The considerations put forth thus indirectly vindicate the view that rational belief is sensitive to higher-order defeat.

2. Lasonen-Aarnio on Higher-Order Defeat

Lasonen-Aarnio frames her discussion within a *rule-based* picture of epistemic rationality on which rational epistemic agents are characterized by following correct epistemic rules. We can think of epistemic rules as abstract functions from epistemic situations to sets of rationally permitted belief states. For example, one epistemic rule might enjoin me to believe that p whenever I have reliable testimony that p , whereas another epistemic rule might enjoin me to believe that p whenever it looks to me as if p . Needless to say, this picture gives rise to a number of tricky questions: how are epistemic rules to be formulated? What does it take for an epistemic rule to be 'correct'? What does it mean to 'follow' an epistemic rule in the relevant sense? I shall largely dodge such foundational issues in what follows.⁵ For present purposes, we can simply grant that a broadly rule-based framework provides a fruitful way of thinking about epistemic rationality.

In line with Lasonen-Aarnio, let us say that a rule-based theory of epistemic rationality is 'single-tiered' given that following correct epistemic rules is both necessary and sufficient for epistemic rationality:

Single-Tiered Framework: An agent's belief state S is epistemically rational if and only if S is the result of following correct epistemic rules.

Suppose we aim to accommodate higher-order defeat within such a single-tiered framework. We must then ensure that whenever an agent in a rational belief state S receives sufficiently strong evidence that S is irrational, S is defeated by that evidence. That is, as Lasonen-Aarnio observes,⁶ epistemic rules must satisfy the following condition in order to be correct:

Rule Condition: If an agent has sufficiently strong overall evidence that her belief state S is not the result of following correct epistemic rules, then S is not the result of following correct epistemic rules.

⁵ See Paul Boghossian, "Epistemic Rules," *Journal of Philosophy* 105 (2008): 472-500 for an illuminating discussion of various foundational issues concerning epistemic rule-following.

⁶ Lasonen-Aarnio, "Higher-Order Evidence and the Limits of Defeat," 321-22.

This condition ensures that whenever an agent in a rational belief state receives a higher-order defeater, the belief state is no longer rational for the agent. Hence, by requiring that epistemic rules must satisfy the Rule Condition to be correct, we have a systematic way of making sense of higher-order defeat within a single-tiered framework.

However, according to Lasonen-Aarnio,⁷ we shouldn't expect that every correct epistemic rule satisfies the Rule Condition. Consider, for instance, a rule that enjoins me to believe that q if I believe that p & $(p \rightarrow q)$. This rule doesn't satisfy the Rule Condition: it's perfectly possible for my belief state to be the result of correctly applying modus ponens, even if I have strong reason to think that modus ponens is invalid or that I'm unable to apply modus ponens correctly. Yet, Lasonen-Aarnio submits, this shouldn't lead us to conclude that modus ponens is incorrect. Rather, we should maintain that modus ponens is correct, but that one can get misleading evidence to the contrary. Later, in §3, I will suggest that we might reasonably hold that modus ponens is an incorrect epistemic rule in virtue of violating the Rule Condition, but maintain that there is nonetheless something importantly right about modus ponens, which isn't diminished by the possibility of acquiring misleading evidence to the contrary. For now, however, let us proceed on the assumption that epistemic rules need not satisfy the Rule Condition to be correct. This means that we must look beyond a single-tiered framework to make sense of higher-order defeat.

According to two-tiered theories, as Lasonen-Aarnio understands them, following correct epistemic rules is necessary, but not sufficient, for epistemic rationality:

Two-Tiered Framework: An agent's belief state S is epistemically rational if and only if:

- (i) S is the result of following correct epistemic rules; and
- (ii) The agent does not have strong overall evidence that S is not the result of following correct epistemic rules.

If we adopt a two-tiered version of the rule-based framework, we can say that a belief state S is irrational despite resulting from following correct epistemic rules. For example, my belief state might be the result of following modus ponens, yet be irrational because I have strong reasons to think that modus ponens is invalid or that I'm unable to correctly apply modus ponens. And John's belief that

⁷ Lasonen-Aarnio, "Higher-Order Evidence and the Limits of Defeat," 322-24.

he is better than average at driving might be the result of following correct inductive rules, yet be irrational because he has strong reasons to think that his belief is the result of an irrational bias. As such, two-tiered theories give us the resources to make sense of higher-order defeat without assuming that epistemic rules must satisfy the Rule Condition to be correct.

But according to Lasonen-Aarnio,⁸ two-tiered theories have a serious drawback, namely that they generate *epistemic dilemmas*: situations where an agent ought to adopt incompatible doxastic attitudes (for instance, believing and suspending judgment about the same proposition). Her reasoning goes as follows: let R be a correct epistemic rule, which enjoins me to believe that p in my current epistemic situation. Assuming that R violates the Rule Condition, it must be possible for me to receive a higher-order defeater in light of which I should give up my belief in p , although R still enjoins me to believe that p . Hence, there must exist another correct epistemic rule R' , which enjoins me to suspend judgment about p upon having received the higher-order defeater. As a result, the rules R and R' end up giving me conflicting recommendations: R enjoins me to believe that p , whereas R' enjoins me to suspend judgment about p . This leads Lasonen-Aarnio to conclude that “[i]f correct rules are accompanied by oughts, it looks like [I] ought to believe that p , and [I] ought to suspend judgment in p .”⁹ In other words, it looks like I face an epistemic dilemma.¹⁰

These are the bare bones of Lasonen-Aarnio’s two-piece challenge to the proponent of higher-order defeat: (i) higher-order defeat seems to call out for a two-tiered theory of epistemic rationality, and (ii) two-tiered theories seem to generate epistemic dilemmas. In her subsequent discussion of this challenge, Lasonen-Aarnio considers and rejects various strategies for meeting the challenge, either by embracing the possibility of epistemic dilemmas, or by trying to specify a version of the rule-based framework (either single-tiered or two-tiered) that doesn’t generate epistemic dilemmas after all. I will return to this part of Lasonen-

⁸ Lasonen-Aarnio, “Higher-Order Evidence and the Limits of Defeat,” 328–30.

⁹ Lasonen-Aarnio, “Higher-Order Evidence and the Limits of Defeat,” 329.

¹⁰ Lasonen-Aarnio, “Higher-Order Evidence and the Limits of Defeat,” 329–30 also argues that cases of conflicting recommendations by correct epistemic rules can arise even if epistemic rules are formulated in terms of what one is permitted to believe, rather than in terms of what one is required to believe. But for present purposes, we can simply grant that epistemic rules specify what is required rather than what is permitted.

Aarnio's discussion in due course, but for now it suffices to have the basic challenge on the table.

3. Higher-Order Defeat in a Single-Tiered Framework

This section concerns the first part of Lasonen-Aarnio's challenge. I want to defend two claims, a negative and a positive. Negatively, I will argue that Lasonen-Aarnio's case against the Rule Condition is ill-founded. Positively, I will use the Rule Condition to outline and motivate a way of accommodating higher-order defeat within a single-tiered framework.

First, the negative point. Recall that Lasonen-Aarnio's case against the Rule Condition was based on the observation that there are seemingly correct epistemic rules, such as *modus ponens*, that nevertheless fail to satisfy the Rule Condition. I think there is something right and something wrong about this observation, which can be brought out by distinguishing two different conceptions of what it means for an epistemic rule to be 'correct.' On one understanding, an epistemic rule is correct to the extent that it is truth-conducive or truth-preserving, where truth-conduciveness is a matter of generating true beliefs, and truth-preservation is a matter of preserving true beliefs.¹¹ On this understanding, *modus ponens* is clearly correct in virtue of being deductively valid. But various non-deductive (e.g. inductive and abductive) rules are presumably also correct in this sense, despite falling short of perfect truth-preservation. On another understanding, an epistemic rule is correct to the extent that it is rational for agents to follow it across all—or at least a sufficiently wide range of—epistemic situations. On this understanding, it is much less clear that *modus ponens* is correct, precisely because one can have strong reason think that *modus ponens* is invalid or that one is unable to apply *modus ponens* correctly. Hence, whether *modus ponens* is a correct epistemic rule may well depend on the notion of correctness we have in mind.

In the present context, since we are considering what it's rational for agents to believe in different epistemic situations, we presumably want to interpret the Rule Condition as a condition on what it takes for an epistemic rule to be correct in the latter sense. Accordingly, when we say that *modus ponens* is incorrect in virtue of violating the Rule Condition, we are not saying that *modus ponens* is anything less than perfectly truth-preserving. Rather, we are saying that it isn't

¹¹ This distinction is equivalent (or at least very similar in spirit) to the distinction between 'conditional' and 'unconditional' reliability in Alvin I. Goldman, "What is Justified Belief?," in *Epistemology. An Anthology*, eds. Ernest Sosa and Jaegwon Kim (Oxford: Blackwell, 1979).

always rational to follow *modus ponens*, since one can have strong reason to think that *modus ponens* is invalid or that one is unable to apply *modus ponens* correctly. And this doesn't strike me as an implausible thing to say, especially if we consider logical rules of a more complex nature than *modus ponens*.

Someone might worry that we end up with a circular account of epistemic rationality, if we spell out what it means for an epistemic rule to be correct in terms of whether it is rational to follow the rule across a sufficiently wide range of circumstances. But there seems to be a natural way of avoiding such circularity by holding that it's rational for an agent to follow a given epistemic rule provided that the agent has (sufficiently strong) overall reason to think that doing so is (sufficiently) conducive to the formation of true beliefs. This allows us to avoid circularity while being faithful to the initial motivation for saying that it can be irrational for an agent to follow an epistemic rule, even if the rule is in fact perfectly truth-conducive or truth-preserving.

Another potential worry about the distinction between a rule's being truth-conducive (or truth-preserving) and its being rational to follow is that it obscures the connection between rationality and truth. If rational belief is supposed to aim at truth, how could it be irrational to follow a rule that is perfectly truth-conducive or truth-preserving? I think we can answer this worry by drawing an analogy between propositions and rules. Just as it can be irrational to believe a true proposition (if one has strong reason to think that the proposition is false), so it can be irrational to follow a truth-conducive or truth-preserving rule (if one has strong reason to think that following the rule isn't conducive to the formation of true beliefs). And just as we can nevertheless aim to believe only what is true, we can nevertheless aim to follow only rules that are truth-conducive or truth-preserving. As such, the distinction between a rule's being truth-conducive (or truth-preserving) and its being rational to follow seems compatible with there being an important normative connection between the two. Obviously, the connection will be less direct than saying that a rule is rational to follow if and only if it is truth-conducive (or truth-preserving). But this seems no less plausible than denying that a proposition is rational to believe if and only if it is true. Hence, the claim that epistemic rules must satisfy the Rule Condition to be correct (in the relevant sense of 'correct') strikes me as quite plausible, or at least not as implausible as Lasonen-Aarnio seems to think.

Next, the positive point. Even if the above considerations are basically correct, it remains to be seen whether there is a satisfactory way of ensuring that

correct epistemic rules satisfy the Rule Condition. In passing Lasonen-Aarnio considers the possibility of amending epistemic rules with provisos for different higher-order defeaters. However, she takes it to be “at least a *prima facie* challenge” to show how an epistemic rule could have provisos “for *all possible* higher-order defeaters built into it.”¹² Here I want to suggest a way of meeting this challenge. As a first step, consider the following qualification of modus ponens:

Modus Ponens (provisional): Follow modus ponens, unless:

- (i) Some correct epistemic rule enjoins you to believe that modus ponens is incorrect; or
- (ii) Some correct epistemic rule enjoins you to believe that you are likely to misapply modus ponens.

This rule features two provisos, which correspond to the two central types of higher-order defeat discussed by Lasonen-Aarnio: (i) higher-order defeat due to evidence that a given epistemic rule is incorrect, and (ii) higher-order defeat due to evidence that one is incapable of applying an epistemic rule correctly. Perhaps some higher-order defeaters do not fall neatly within these two categories. If so, we might have to amend modus ponens with additional provisos. But the provisos included here cover at least the most commonly discussed cases of higher-order defeat in the recent literature.

If all correct epistemic rules come with provisos akin to those in Modus Ponens (provisional), we have the resources to accommodate higher-order defeat within a single-tiered framework. Consider a structurally similar case to Self-Enhancement Bias, adapted from Christensen:¹³

Modus Ponens on Drugs: Suzy rationally believes that $p \ \& \ (p \rightarrow q)$, and she forms a belief in q as a result of correctly applying modus ponens. But she is now told from a reliable source that she has been given a reason-distorting drug that subtly, but significantly, impairs her ability to perform even simple deductive inferences. As a matter of fact, Suzy has not been drugged.

Assuming that this is a case of higher-order defeat, Suzy is rationally permitted to believe that q *before* being told that she has been drugged, but she isn’t rationally permitted to believe that q *after* being so told. And the proposed single-tiered account yields precisely this result, since the second proviso in Modus

¹² Lasonen-Aarnio, “Higher-Order Evidence and the Limits of Defeat,” 324.

¹³ David Christensen, “Disagreement, Drugs, Etc.: From Accuracy to Akrasia,” *Episteme* 13 (2016): 401.

Ponens (provisional) is met *after*, but not *before*, Suzy is told about the drug. (The same would obviously hold if Suzy's higher-order defeater had instead consisted of evidence indicating that modus ponens is invalid.)

However, while epistemic rules like Modus Ponens (provisional) allow us to make sense of *ordinary* cases of higher-order defeat—like Modus Ponens on Drugs and Self-Enhancement Bias—they do not allow us to make sense of *all possible* cases of higher-order defeat. For as Lasonen-Aarnio¹⁴ points out, Suzy could in principle get strong reason to think that Modus Ponens (provisional) is incorrect or that she is unable to apply it correctly. Hence, if we want to accommodate such non-standard cases of higher-order defeat, we need to amend modus ponens with a second layer of provisos:

Modus Ponens (doubly provisional): Follow Modus Ponens (provisional), unless:

- (i) Some correct epistemic rule enjoins you to believe that Modus Ponens (provisional) is incorrect; or
- (ii) Some correct epistemic rule enjoins you to believe that you are likely to misapply Modus Ponens (provisional).

Continuing this path, we run into an infinite regress: a third layer of provisos is needed to accommodate a case where Suzy gets strong reason to think that Modus Ponens (doubly provisional) is incorrect or that she is unable to apply it correctly, and so on. Thus, if we want to accommodate all possible cases of higher-order defeat, it looks like correct epistemic rules must feature infinite layers of provisos for different kinds of higher-order defeat.

What should we make of the resulting view? By way of comparison, consider a different single-tiered view, which Lasonen-Aarnio discusses under the label 'The Über-Rule View'.¹⁵ On this view, a belief state is epistemically rational if and only if it is the result of following an overarching epistemic rule—the Über-rule—which determines, for any given epistemic situation, which belief states are rationally permitted in that situation. Accordingly, higher-order defeat is to be made sense of by reverse-engineering the content of the Über-rule in such a way that it delivers the desired verdicts in cases like Self-Enhancement Bias.

While Lasonen-Aarnio raises a number of different worries about The Über-Rule View, perhaps the most serious problem is that an extensionally adequate Über-rule must presumably be a highly complex construct of gerrymandering,

¹⁴ Lasonen-Aarnio, "Higher-Order Evidence and the Limits of Defeat," 323.

¹⁵ Lasonen-Aarnio, "Higher-Order Evidence and the Limits of Defeat," sect. 4.

which can hardly be expressed in “any finite, informative way.”¹⁶ Consequently, the worry goes, the Über-rule can hardly offer the kind of epistemic guidance that we should expect correct epistemic rules to be able to offer.¹⁷

Let us straightaway grant this criticism of the Über-rule view. Does the single-tiered view outlined above fall prey to a similar criticism? It seems not. Despite the addition of infinite layers of provisos, the view is nevertheless able to offer genuine guidance. Consider Modus Ponens on Drugs: we previously saw that Suzy’s initial belief that q is deemed rational, since none of the provisos in Modus Ponens (provisional) are initially met. But when Suzy receives the higher-order defeater, she may no longer believe that q , because the second proviso in Modus Ponens (provisional) is now met. Hence, it looks like Suzy has all the guidance she needs to decide which doxastic attitude she should adopt towards q before and after receiving the higher-order defeater. And the same obviously goes for other ordinary cases of higher-order defeat such as Self-Enhancement Bias.

What about non-standard cases of higher-order defeat? Suppose Suzy gets a second higher-order defeater, which consists of evidence that Modus Ponens (provisional) is incorrect. Here things become a bit more complicated. On the proposed view, Suzy should resort to Modus Ponens (doubly provisional), because one of the provisos in Modus Ponens (provisional) is met. But which doxastic attitude does this rule enjoin Suzy to adopt towards q ? It seems to depend on the exact content of the second higher-order defeater. Suppose the second higher-order defeater indicates that Modus Ponens (provisional) is incorrect, without indicating *why* it is incorrect. We can then imagine how Suzy might entertain different possible explanans of why Modus Ponens (provisional) is incorrect. One possible explanation is that it is always reasonable to follow modus ponens, contrary to what Modus Ponens (provisional) says. If so, Suzy is presumably enjoined to believe that q , since this is what modus ponens enjoins. But another possible explanation is that Modus Ponens (provisional) is incorrect in virtue of failing to take into account every type of situation in which it is unreasonable to follow modus ponens. If so, Suzy is presumably enjoined to suspend judgment about q , since her current situation still counts as one in which it is unreasonable to follow modus ponens. Thus, it looks like Suzy faces a problem of underdetermination about which doxastic attitude she is enjoined to adopt towards q .

¹⁶ Lasonen-Aarnio, “Higher-Order Evidence and the Limits of Defeat,” 332.

¹⁷ Lasonen-Aarnio, “Higher-Order Evidence and the Limits of Defeat,” 333.

Is this sort of underdetermination a vice or a virtue of the proposed view? Someone who thinks that there must always be a determinate answer to questions about which doxastic attitudes we should adopt towards different propositions might want to conclude that the proposed single-tiered theory fails to offer enough guidance in situations like Suzy's. However, it isn't clear to me that there *is* a determinate answer to the question of which doxastic attitude Suzy should adopt towards q . Suzy's situation might just be a genuine case of underdetermination concerning which doxastic attitude she should adopt. If so, we should consider the underdetermination property of the proposed view as a feature, not a bug.

In any case, regardless of how we settle the issue about underdetermination, the proposed view at least seems well placed to offer genuine guidance in ordinary cases of higher-order defeat like Modus Ponens on Drugs and Self-Enhancement Bias. As such, I see no reason to be pessimistic about the prospects for accommodating higher-order defeat within a single-tiered framework.

4. Higher-Order Defeat in a Two-Tiered Framework

This section concerns the second part of Lasonen-Aarnio's challenge. My aim is to argue that even if the single-tiered view outlined in the previous section cannot be made to work, we can still make sense of higher-order defeat within a two-tiered framework without committing ourselves to the possibility of epistemic dilemmas.

Recall Lasonen-Aarnio's basic reasons for thinking that two-tiered theories tend to generate epistemic dilemmas: (i) correct epistemic rules can give conflicting recommendations in a two-tiered framework, and (ii) correct epistemic rules are "accompanied by oughts."¹⁸ My strategy will be to grant the former claim, but deny the latter: correct epistemic rules may well deliver conflicting recommendations in a two-tiered framework, but we should not expect such conflicting recommendations to constitute genuine epistemic dilemmas.

Let us begin with a simple observation: epistemic rules play a very different role in a two-tiered framework than in a single-tiered framework. Suppose a rule R enjoins an agent to believe that p in a given situation. On a single-tiered view, R 's recommendation will invariably constitute a rational requirement on part of the agent to believe that p , since single-tiered theories work with a straightforward 'one-to-one' correspondence between requirements of rationality and recommendations of correct epistemic rules. Accordingly, conflicting

¹⁸ Lasonen-Aarnio, "Higher-Order Evidence and the Limits of Defeat," 329.

recommendations by correct epistemic rules always generate epistemic dilemmas within a single-tiered framework. By contrast, in a two-tiered framework, R 's recommendation need not constitute a rational requirement to believe that p , since two-tiered theories do not feature a one-to-one correspondence between requirements of rationality and recommendations of correct epistemic rules. Accordingly, conflicting recommendations by correct epistemic rules need not generate epistemic dilemmas within a two-tiered framework.

This observation already suffices to mitigate, at least to some degree, the second part of Lasonen-Aarnio's challenge. Suppose I'm enjoined to believe that p by a correct epistemic rule R . Assuming, as we do, that R need not satisfy the Rule Condition, it must be possible for me to acquire a higher-order defeater in light of which I should give up my belief that p , although R still enjoins me to believe that p . Hence, there must exist another correct epistemic rule R' , which enjoins me to suspend judgment about p upon having received the higher-order defeater. As a result, the rules R and R' end up giving me conflicting recommendations: R enjoins me to believe that p , whereas R' enjoins me to suspend judgment about p . This is what led Lasonen-Aarnio to conclude that two-tiered theories tend to generate epistemic dilemmas. However, since R and R' make their conflicting recommendations in a situation where I have received a higher-order defeater, proponents of the two-tiered view will presumably want to say that R 's recommendation no longer constitutes a rational requirement to believe that p , in which case I do not face an epistemic dilemma. After all, it was the ability to accommodate such cases of higher-order defeat that motivated the move to a two-tiered framework to begin with. Hence, we should not expect two-tiered theories to generate epistemic dilemmas in cases of higher-order defeat.

Obviously, this does not yet show that there is a satisfying way of filling in the details of a two-tiered theory, which ensures that conflicting recommendations by correct epistemic rules never constitute epistemic dilemmas. In particular, it does not fend off Lasonen-Aarnio's criticism of the specific two-tiered view that she calls 'The Hierarchy View.'¹⁹ On this view, a belief state is epistemically rational if and only if it is the result of following a correct epistemic rule, which isn't overridden by a correct epistemic rule higher up the hierarchy. The hierarchy of rules is then determined by a meta-rule, which assigns an ordering of epistemic rules to each epistemic situation. Hence, given the right meta-rule, The Hierarchy View might help avoid epistemic dilemmas in cases of higher-order defeat: if R'

¹⁹ Lasonen-Aarnio, "Higher-Order Evidence and the Limits of Defeat," sect. 6.

overrides R , I will be required to suspend judgment about p upon having acquired the higher-order defeater, but I will not be required to believe that p .

However, just as Lasonen-Aarnio deems The Über-Rule View inapt to accommodate higher-order defeat within a single-tiered framework, she deems The Hierarchy View inapt to accommodate higher-order defeat within a two-tiered framework. Here the main worry concerns what happens if one gets strong reason to think that the meta-rule is incorrect or that one is unable to apply it correctly. According to Lasonen-Aarnio, if we want to accommodate such cases of higher-order defeat, we must introduce a set of meta-rules together with a ‘meta-meta-rule,’ which gives us a ranking of the meta-rules, for each epistemic situation. But this strategy quickly leads to an infinite regress due to the further possibility of acquiring reasons to think that the meta-meta-rule is incorrect or that one is unable to apply it correctly. Hence, according to Lasonen-Aarnio, we face the same worry that led us to reject The Über-Rule View: if The Hierarchy View must feature infinitely many orders of meta-rules, it becomes unclear how the view can offer genuine guidance.²⁰

This guidance-worry seems to have less force against The Hierarchy View than against The Über-Rule View. Of course, there is a sense in which neither view is “finitely expressible,” but The Hierarchy View nevertheless seems informative in a way that The Über-Rule View is not. In fact, The Hierarchy View seems informative in much the same way as the single-tiered view proposed in the previous section. Consider Modus Ponens on Drugs: on The Hierarchy View, Suzy’s belief that q is initially rational, since it is the result of following a correct epistemic rule, call it R , which is not overridden by any other correct epistemic rule. But when receiving the higher-order defeater, Suzy is no longer rationally permitted to believe that q , since R ’s recommendation is now overridden by another rule, call it R' , which recommends her to suspend judgment about p . Hence, it looks like Suzy has all the guidance she needs to determine which doxastic attitude to adopt towards q before and after receiving the higher-order defeater. Obviously, The Hierarchy View allows for the possibility that the recommendation by R' is itself overridden at some later stage, in which case it might become less clear what The Hierarchy View recommends (for reasons explained in the previous section). But at the very least, the Hierarchy seems able to offer genuine guidance in ordinary cases of higher-order defeat like Modus Ponens on Drugs and Self-Enhancement Bias. Hence, I see no reason to think that the kind of infinity

²⁰ Lasonen-Aarnio, “Higher-Order Evidence and the Limits of Defeat,” 340–41.

featuring in The Hierarchy View is detrimental to the view. But even if it turns out that The Hierarchy View cannot ultimately be made to work, it would be premature to conclude that no two-tiered theory can be made to work, unless there are general grounds for expecting that two-tiered theories tend to generate epistemic dilemmas in cases of higher-order defeat. And I take the foregoing considerations to show that such general grounds are lacking.

THE WARRANT ACCOUNT AND THE PROMINENCE OF ‘KNOW’

Jacques-Henri VOLLET

ABSTRACT: Many philosophers agree that there is an epistemic norm governing action. However, they disagree on what this norm is. It has been observed that the word ‘know’ is prominent in ordinary epistemic evaluations of actions. Any opponent of the knowledge norm must provide an explanation of this fact. Gerken has recently proposed the most developed explanation. It invokes the hypothesis that, in normal contexts, knowledge-level warrant is frequently necessary and very frequently sufficient (Normal Coincidence), so that knowledge-based assessments would be a good heuristic for practical reasoning and epistemic evaluations of action. In this paper, I raise three problems for this approach. First, I argue that Normal Coincidence is ad hoc: it relies on an unsupported frequency hypothesis that we should expect to be false given the warrant account that Gerken also endorses. Second, I argue that, in any case, Normal Coincidence is insufficient to support the hypothesis that knowledge-based evaluation of action constitutes a good heuristic. Third, I consider three other hypotheses close to Normal Coincidence apparently more likely to support the heuristic hypothesis, but I argue that they seem even more ad hoc than Normal Coincidence.

KEYWORDS: norm, epistemic, action, knowledge, practical reasoning, heuristic

1. Introductory Remarks

We often say things like

- “You *knew* that p (e.g., that he is allergic to peanuts). Why didn’t you tell me?”
- “She did well to do A (e.g., to check), for she *had insufficient evidence*.”

This suggests that there is an epistemic norm governing action.¹ This norm is supposed to tell us which epistemic condition must be satisfied to rationally (or

¹ Although see Davide Fassio, “Is There an Epistemic Norm of Practical Reasoning,” *Philosophical Studies* 174, 9 (2017), for the idea that such epistemic evaluations concern regulation conditions rather than a (sui generis) epistemic norm of practical reasoning. On the distinction between norms and regulation conditions, see also Pascal Engel, “Belief and Normativity,” *Disputatio* 2, 23 (2007) and “In What Sense Is Knowledge the Norm of Assertion?” *Grazer Philosophische Studien* 77, 1 (2008).

appropriately) use a certain consideration as a reason for action.² There is no consensus on what the relevant epistemic condition is.³ Still, everybody accepts that ‘know’ is a prominent term of ordinary epistemic assessments of action and that this fact constitutes some evidence in favour of the knowledge account. Thus, any opponent of the knowledge account must offer a plausible explanation of it.

Mikkel Gerken, an opponent of the knowledge norm, has recently offered the most promising and developed explanation of the required sort. This explanation is based on the hypothesis that, in normal contexts, knowledge-level warrant is frequently necessary and very frequently sufficient for action (Normal

² The supposed norm has received different formulations. Some philosophers formulate the norm in terms of a condition for the rationality of action, others in terms of a condition for the appropriateness, warrant or permissibility of action. What the norm is supposed to govern is also expressed in various ways, for example in terms of “acting as if p,” “acting for the reason that p,” “treating p as a reason,” “acting on p,” “using p as a premise in one’s practical reasoning and action” etc. Although these differences in formulation may matter (see Ram Neta, “Treating Something as a Reason for Action,” *Noûs* 43, 4 (2009): 684–686; Mikkel Gerken, “Warrant and Action,” *Synthese* 178, 3 (2011): 535n9) they won’t concern us here.

³ Some philosophers say that it is knowledge. See e.g. John Hawthorne, *Knowledge and Lotteries* (Oxford: Oxford University Press, 2004); Jason Stanley, *Knowledge and Practical Interests* (Oxford: Oxford University Press, 2005); Timothy Williamson, “Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge,” *Philosophical Quarterly* 55, 219 (2005); John Hawthorne and Jason Stanley, “Knowledge and Action,” *Journal of Philosophy* 105, 10 (2008). Other theorists argue that it is knowledge-level justification. See for example Jeremy Fantl and Matthew McGrath, *Knowledge in an Uncertain World* (Oxford: Oxford University Press, 2009); Clayton Littlejohn, “Must We Act Only on What We Know?” *Journal of Philosophy* 106, 8 (2009); Clayton Littlejohn, *Justification and the Truth-connection* (Cambridge: Cambridge University Press, 2012). According to Neta, “Treating Something,” it is justification to believe that one knows. Some other writers think that the epistemic norm provides an epistemic condition sensitive to the practical context. See, among others, Jessica Brown, “Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning,” *Noûs* 42, 2 (2008); Baron Reed, “A Defense of Stable Invariantism,” *Noûs* 44, 2 (2010); Gerken, “Warrant and Action,” Mikkel Gerken, “The Roles of Knowledge Ascriptions in Epistemic Assessment,” *European Journal of Philosophy* 23, 1 (2015); Mikkel Gerken, *On Folk Epistemology. How We Think and Talk about Knowledge* (Oxford: Oxford University Press, 2017); Dustin Locke, “Practical Certainty,” *Philosophy and Phenomenological Research* 90, 1 (2015). Finally, there are philosophers who embrace a pluralist view according to which there are various epistemic norms. See in particular Janet Levin, “Assertion, Practical Reason, and Pragmatic Theories of Knowledge,” *Philosophy and Phenomenological Research* 76, 2 (2008).

Coincidence), so that knowledge-based evaluation constitutes a good heuristic for evaluating action and practical reasoning.⁴

In this paper, I raise three problems for this approach. First, I argue that Normal Coincidence (NC) is ad hoc. It relies on an unsupported frequency assumption that we should expect to be false given the warrant account that Gerken also endorses.⁵ Second, I show that, in any case, NC is insufficient to support the hypothesis that knowledge-based evaluation constitutes a good heuristic, for NC does not exclude that knowledge-level warrant is not frequently necessary and sufficient for appropriate action, nor that knowledge-level warrant is not sufficiently close to the degree of warrant which is (very) frequently necessary and sufficient for action. Third, I consider three alternative frequency hypotheses close to NC but apparently more fit to support the hypothesis that knowledge-based assessment of action constitutes a good heuristic. I show that they are even more ad hoc than NC.

My plan is as follows. In section 2, I clarify the claim that knowledge-based assessments are prominent in ordinary epistemic assessments of action by distinguishing three ways in which that is the case and how NC is supposed to explain this. In section 3, I explain why NC is ad hoc and insufficient and why the three alternatives are even more ad hoc. In section 4, I anticipate a possible reply appealing to the distinction between normal and abnormal situations and I show that it is unsatisfactory.

2. Explaining the Prominence of ‘Know’ in Epistemic Assessments of Actions

2.1 Three Ways in Which ‘Know’ Is Prominent

There are at least three ways in which the word ‘know’ is prominent in epistemic assessments of action. First, ‘know’ seems to be the *default* word of epistemic assessments of action. To see this, consider the following case, from Hawthorne and Stanley:

RESTAURANT. Hannah and Sarah are trying to find a restaurant, at which they have time-limited reservations. Instead of asking someone for directions, Hannah goes on her hunch that the restaurant is down a street on the left. After walking for some amount of time, it becomes quite clear that they went down the wrong

⁴ Gerken, “The Roles of Knowledge Ascriptions” and *On Folk Epistemology*.

⁵ Gerken, “Warrant and Action” and *On Folk Epistemology*.

street.⁶

In this situation, it may be granted that Hannah's action is wrong because a mere hunch that *p* is insufficient to act on *p*. However, as Hawthorne and Stanley emphasize,

[a] natural way for Sarah to point out that Hannah made the wrong decision is to say, "You shouldn't have gone down this street, since you didn't know that the restaurant was here."⁷

Here, we may think that the use of 'know' is natural in the sense that, by default (in the absence of further information), it is appropriate to use 'know' rather than, e.g., 'justified' (or 'rational') belief. The fact that 'know' is the appropriate default term of evaluation constitutes a first way in which 'know' is prominent. Call this 'the default claim.'

Default claim: 'know' is the default term of epistemic assessments of action.

It is also remarkable that 'know' is the most frequently used term in our epistemic assessments of action.⁸ This fact constitutes a second way in which 'know' is prominent. Call this 'the frequency claim.'

Frequency claim: 'know' is the most frequently used term for epistemic assessment of action

Finally, 'knowledge' questions and ascriptions always sound relevant and appropriate when it comes to making a practical decision. As Reed (an opponent of the knowledge norm) writes:

[Y]ou are trying to decide whether to check if the train stops in Foxboro because it is extremely important that you get there as quickly as possible. You have not yet decided whether it is rational to check if the train makes that stop. One of the relevant factors in your decision would presumably be an answer to the question, do you know the train will stop there?⁹

⁶ Hawthorne and Stanley, "Knowledge and Action," 571.

⁷ Hawthorne and Stanley, "Knowledge and Action," 571.

⁸ As far as I know, no systematic study showing that 'know' is prominent in epistemic appraisal of action has ever been conducted, but 'know' is one of the ten words most used (see Gerken, *On Folk Epistemology*, 15) and, to my knowledge, all philosophers agree with the frequency claim.

⁹ Reed, "A Defense of Stable Invariantism," 232.

We should also note that a positive answer to a 'knowledge' question is always decisive for action, even when the stakes are high.¹⁰ For example, suppose you ask Jill whether you should take a certain train. Suppose she tells you "I know that this train will stop there. You should take it!" It would be quite odd for you to reply "But the question is not whether you know!" By saying "I know that this train will stop there," Jill provides you with appropriate information with respect to whether you should take this train and act on the proposition that the train will stop there. The third way in which 'know' is prominent has to do with the role that 'know' plays in appropriate questions and recommendations when it comes to action. Call this the 'guidance claim':

Guidance claim: 'know' is always appropriate in questions and recommendations about what to do.

The fact that 'know' is prominent in epistemic assessments of action constrains any satisfactory account to provide a plausible explanation of this phenomenon.¹¹ Of course, the offered explanation has to be consistent with the assumed epistemic norm. This constitutes a challenge for philosophers rejecting the knowledge norm.¹²

2.2. Explaining the Prominence of 'Know' on the Basis of Normal Coincidence

Gerken is a classical invariantist about knowledge. According to him, the epistemic standards for knowledge (or 'know') do not shift with the practical context of the subject (or of the attributor). Regarding the epistemic norm of action, Gerken

¹⁰ See Jeremy Fantl and Matthew McGrath, "On Pragmatic Encroachment in Epistemology," *Philosophy and Phenomenological Research* 75, 3 (2007): 562.

¹¹ Gerken explicitly understands the prominency claim as a frequency claim given normal situations:

Prominence of 'Knowledge' In normal cases of epistemic assessment of action and assertion, ordinary speakers frequently use the term 'knowledge' and its cognates (Gerken, *On Folk Epistemology*, 17).

But, as I explain below, he also grants that 'know' is the default term (and that the concept *knowledge* is the default concept) of epistemic assessments, and that 'knowledge' ascriptions have a "directive force" for action.

¹² Proponents of the knowledge account face a different challenge, though, that of explaining why we do not always use 'know' to assess actions. See Brown, "Subject-sensitive Invariantism" and Gerken, *On Folk Epistemology*, 18.

embraces the warrant account, according to which the epistemic norm of action provides different epistemic standards associated with different deliberative contexts, where a deliberative context is determined by a set of practical factors. More precisely:

Warrant Account (WA). In the deliberative context, DC, S meets the epistemic conditions on rational use of (her belief that) p as a premise in practical reasoning or of (her belief that) p as a reason for acting (if and) only if S is warranted in believing that p to a degree that is adequate relative to DC.¹³

The notion of warrant that Gerken invokes is non-factive. Therefore, on this view, knowledge is never necessary to warrant action. The notion of warrant is also gradable, and the practical factors that determine the degree of warrant which is adequate include in particular the urgency of action, the availability of evidence and other options, the stakes, and the social role of the agent.

In order to explain the prominence of ‘know,’ Gerken appeals to the frequency assumption that, in normal cases, knowledge-level warrant is frequently necessary and very frequently sufficient:

Normal Coincidence (NC). In normal cases of epistemic assessment, the degree of warrant necessary for S’s knowing that p is frequently necessary and very frequently sufficient for the epistemic permissibility of S’s acting on (the belief that) p.¹⁴

It’s important to note that by ‘normal,’ Gerken does not mean the most frequent for an individual. As we will see, according to him, a surgeon may frequently face an abnormal situation (see the discussion in section 4).

How is NC supposed to explain the prominence of ‘know’? The idea seems to be that if NC is correct then, in normal situations, it is frequently right, or at least frequently sufficiently close to being right, to use this term. This would explain the frequency claim.¹⁵

¹³ Gerken, “Warrant and Action,” 530.

¹⁴ Gerken, “The Roles of Knowledge Ascriptions” and *On Folk Epistemology*, 143.

¹⁵ Note, though, that ‘know’ is a factive term. If the warrant required for action is never factive, it is odd that we use a factive term. Gerken acknowledges this point and appeals to the idea that ‘know’ is learnt before other epistemic terms (see Gerken, “The Roles of Knowledge Ascriptions”). Perhaps Gerken could also say that, typically, what is knowledge-level warranted is believed and true, and hence, known, so that there is a rough coincidence between what is knowledge-level warranted and what is known. See also the discussion of Douven below.

NC could also explain the default claim. Again, assume that, given NC, knowledge-level warrant is frequently necessary and sufficient in normal situations, or at least sufficiently close to the required warrant for action. Then, epistemic assessments of action in terms of knowledge could be seen as relying on a good heuristic process appealing to knowledge as a good “cognitive proxy.”¹⁶ This strategy would manage to “provide a good trade-off between accuracy and communicative effectiveness.”¹⁷ This approach has it that we use the concept *knowledge* by default in our epistemic assessments. Indeed, default conceptualisations must be informative and easy to process. We may think that conceptualisations in terms of *know* are better placed than others in this respect. More fine-grained conceptualisations (in terms of seeing, remembering, etc.) would be more costly, involving the multiplication of more complex heuristics. In contrast, by using the concept *knowledge* in our epistemic assessments, we could use this concept as a middle-term in heuristic inferences like, for example, knowledge-output heuristics (e.g., from “S said that p,” to “S knows that p”) and knowledge-input heuristics (e.g., from “S knows that p” and “S knows that q,” to “S knows that r”). More coarse-grained conceptualisations (e.g., in terms of belief or reliable belief), on the other hand, would involve too great an information loss.¹⁸ If the concept *know* is the concept we use by default to make epistemic assessments of action, it’s not surprising that we use the word ‘know’ by default in these assessments (since ‘know’ expresses the concept *know*).

Let us now turn to the guiding claim according to which ‘know’ is prominent in the sense that ‘know(s)’ questions and ascriptions are always appropriate and relevant when it comes to making a practical decision. To explain this phenomenon, Gerken invokes a supposed pragmatic feature of ‘knowledge’ ascriptions.¹⁹ The starting idea is that an assertion that p can have different illocutionary effects, and hence, fulfil different communicative functions. A potential illocutionary effect of assertions is that of recommending a certain course of action. For example, by asserting “The window is open,” you can perform the indirect speech act of commanding someone to close the window. Likewise, ‘knowledge’ ascriptions are assertions. They can serve the communicative function of directing action. As with conversational implicatures, the directive force

¹⁶ Gerken, *On Folk Epistemology*, 149.

¹⁷ Gerken, *On Folk Epistemology*, 190.

¹⁸ Gerken, *On Folk Epistemology*, 101-104.

¹⁹ Gerken, *On Folk Epistemology*, ch. 8.

associated with 'knowledge' ascriptions is to some extent context-dependent. We can then distinguish *particularised* implicatures which heavily rely on the particularity of the context and the background assumptions ascribed to the interlocutors, and *generalised* implicatures, which constitute the default interpretation of the speech act. For example, "Can you pass the salt?" generally implicates a request for the salt. This is the default interpretation.

Given that an implicature can depend more or less on the context, Gerken articulates a notion of more or less *regularised* implicature, whereby it is registered how strongly the implicature depends on the particularity of the context. He proposes to say that in contexts where the question is to decide what to do, 'knowledge' ascriptions have a *highly regularised* directive illocutionary effect, such that the default interpretation is a recommendation to act in a certain way.

A possible worry here is that we cannot just assume that there is a highly regularised directive illocutionary effect associated with 'knowledge' ascriptions. As DeRose notes, an explanation postulating that 'knowledge' ascriptions have a pragmatic meaning besides their semantic meaning would be plainly ad hoc.²⁰ However, Gerken's pragmatic explanation partially relies on the claim that the concept *knowledge* is the default concept of epistemic assessments. As he writes:

If the concept *knowledge* is used by default in intuitive epistemic judgments, it is reasonable to assume that the word 'knowledge' is used in a similar manner as a *communicative heuristic*.²¹

In other words, it is suggested that, in the same way as we use the concept *knowledge* as an heuristic for epistemic assessments, we use the word 'know' as an heuristic to communicate these assessments. If this is correct, it is not surprising that 'knowledge' ascriptions (typically) have the function of recommending or not recommending various actions in contexts where the point is to make a practical decision. In this respect, the hypothesis of a regularised pragmatic effect of 'knowledge' ascriptions is not ad hoc.

In summary, the explanation of the guiding claim requires the assumption that using the concept *knowledge* is a good heuristic for epistemic assessments of action. The plausibility of this assumption is supposed to rely in turn upon the truth of Normal Coincidence. For if the epistemic warrant required for action is, in

²⁰ Keith DeRose, *The Case for Contextualism: Knowledge, Skepticism, and Context*, Vol. 1. (Oxford: Oxford University Press, 2009), 120.

²¹ Gerken, *On Folk Epistemology*, 189.

normal contexts, frequently (much) higher or lower than knowledge-level warrant, the supposed heuristic could not constitute a good trade-off. Thus, although Gerken's proposal invokes default conceptualisations, Normal Coincidence remains the cornerstone of his general approach.

3. Normal Coincidence

It is now time to assess Gerken's proposal. As we have seen, it consists of three essential claims. First, there is WA. Second, there is NC. Third, there is the claim that NC explains the prominence of 'know' by supporting the hypothesis that knowledge-based assessment constitutes a good heuristic for practical reasoning and epistemic evaluation of action.

To begin with, consider NC. NC seems to rely on an unsupported frequency assumption, namely, that knowledge-level warrant is frequently necessary. Indeed, according to Gerken, low stakes cases such as the following do not require knowledge-level warrant:

KICKOFF. S believes that the game has started. But the only basis for her belief is that she vaguely remembers a stranger telling her the time of the kick-off in the bar the night before. But both S and the testifier were tipsy, and the fellow didn't seem all that reliable anyhow.²²

Suppose that S turns on the TV on the basis of his belief that the game has started, and assume that, given that the cost of error is low and the reward in accuracy is high, S's action is warranted. There is no reason to think that this kind of situation is exceptional. There are plenty of cases where we do not have much evidence in favour of a certain proposition, and still act on it. If the warrant account is true, these situations are typically ones in which a significantly low warrant for action is required, presumably a warrant lower than knowledge-level. If we frequently face these kinds of situations, then a warrant lower than knowledge-level is frequently sufficient. Since it's not at all implausible that we very frequently face very low stakes situations, the assumption that knowledge-level warrant is frequently necessary for rational action is at best ill-motivated.

One might want to reply that knowledge-level warrant is not that high. We frequently reach it in ordinary life. But this reply makes the claim that knowledge-level warrant is *very* frequently sufficient less plausible. For, if knowledge-level warrant is low, we should expect it to be *very* frequently *necessary* and merely

²² Gerken, *On Folk Epistemology*, 145.

frequently sufficient. Since Normal Coincidence has it that knowledge-level warrant is very frequently sufficient, it suggests that the standards for knowledge are not that low. But then, why assume that in frequent (normal) ordinary contexts, knowledge-level warrant is necessary? Although these considerations do not constitute a knock-down argument against NC, they suggest that NC is ad hoc, for NC is based on an unsupported frequency assumption.

Further, given the warrant account, we should expect NC to be false. Indeed, according to NC, normal practical situations are roughly uniform regarding the degree of warrant required for appropriate or rational action. But according to WA, there are many parameters which are relevant to determine the deliberative context. They can combine in many different ways. Presumably, each one can affect the epistemic warrant more or less strongly. But, if so, given all the ways in which these parameters can combine and the different degrees to which they can affect the epistemic warrant, it is very surprising that these parameters provide in frequent normal cases more or less the same epistemic condition. This is a second reason to think that NC is ad hoc.

The problem comes from the fact that, on Gerken's view, knowledge-level warrant is supposed to be independent of the practical situation and invariant, whereas the degree of warrant required for action is supposed to be sensitive to many practical factors to various degrees, and hence, very flexible. Given this, NC looks like a miracle.

On this score, compare the strategy deployed by Gerken to the strategy deployed by Douven with respect to the norm of assertion. Douven argues in favour of a rational credibility norm for assertion. According to him, we may explain why 'know' is a prominent term of epistemic assessment of assertions on the basis of the consideration that most of the time, what is rationally credible is known (most of the time, we believe what is rational to believe and our rational beliefs are true and not true by mere luck).²³ On this view, there is a normal coincidence between knowledge and rational credibility because the epistemic requirement for knowledge and rational credibility is the same. But if there is no independent reason to think that normal practical contexts are (roughly) uniform regarding the different parameters that they involve, NC is most unexpected.

Consider now the claim that NC explains the prominence of 'know.' NC explains the prominence of 'know' only if it supports the hypothesis that

²³ Igor Douven, "Assertion, Knowledge, and Rational Credibility," *Philosophical Review* 115, 4 (2006): 469-470.

The Warrant Account and the Prominence of ‘Know’

knowledge-based assessments constitute a good heuristic for practical reasoning and epistemic evaluation of action. This seems to suggest that NC can explain the prominence of ‘know’ only if it implies either that (1) knowledge-level warrant is frequently necessary and sufficient (in normal contexts), or (2) knowledge-level warrant is frequently close enough to the warrant which is frequently necessary and sufficient (in normal contexts).

Consider the first possibility:

Normal Coincidence* (NC*). In normal cases of epistemic assessment, the degree of warrant necessary for S’s knowing that p is frequently necessary and sufficient for the epistemic permissibility of S’s acting on (the belief that) p

The idea is that if knowledge-level warrant were frequently not necessary and sufficient (in normal contexts), then using ‘know’ (and the concept *know*) in ordinary assessments of action would be frequently too strong or too weak. Using ‘know’ (and the concept *know*) would be frequently inaccurate (in normal contexts) and could not constitute a good trade-off.

However, NC does not entail NC*. Suppose that there are seven cases out of ten in which knowledge-level warrant is necessary, so that we may say that knowledge-level warrant is frequently necessary. And suppose that there are nine cases out of ten where knowledge-level warrant is sufficient, so that we may say that knowledge-level warrant is very frequently sufficient. In this situation, it’s possible that only six cases out of ten are cases in which knowledge-level is necessary and sufficient. This may mean that knowledge-level warrant is not frequently necessary and sufficient, even if knowledge-level warrant is frequently necessary and very frequently sufficient.

The following table illustrates the situation (where ‘N’ stands for ‘knowledge-level warrant is necessary’ and ‘S’ stands for ‘knowledge-level warrant is sufficient’):

S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
N	N	N	N	N	N	N			
	S	S	S	S	S	S	S	S	S
	N&S	N&S	N&S	N&S	N&S	N&S			

Of course, it is possible to stipulate afterwards a threshold for “frequent” and “very frequent” such that in the case under consideration knowledge-level warrant is frequently necessary and sufficient. For example, we may stipulate that “six cases out of ten” counts as “frequent.”

However, this does not get us very far, for if we postulate that “six cases out of ten” counts as “frequent,” then the following table is possible given NC:

S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
N	N	N	N	N	N				
	S	S	S	S	S	S	S	S	S
	N&S	N&S	N&S	N&S	N&S				

In this table, knowledge-level warrant is frequently necessary (by assumption) and very frequently sufficient, but it seems false that knowledge-level warrant is frequently necessary and sufficient, for there are only five cases of that sort. If one wants to maintain afterwards that “five cases out of ten” can count as “frequent,” we can repeat the operation.

Second, we should note that the lower the assumed threshold for “frequent,” the less plausible is the claim that using the concept *knowledge* and the word ‘know’ constitute good heuristics for practical reasoning and epistemic assessments of action. For example, we may accept that “six cases out of ten” counts as “frequent,” but then it becomes difficult to argue on this basis that using the concept *knowledge* and the word ‘know’ constitutes a good trade-off. Gerken’s approach does not fit well with the idea that the threshold for “frequent” could be that low.

In any case, it’s clear that the more cases we postulate in which knowledge-level warrant is necessary and sufficient, the more ad hoc and implausible the proposal is. Indeed, assuming WA and NC, we have reason to expect that cases in which knowledge-level warrant is necessary and sufficient should be the exception, rather than the rule.

First, suppose (following NC) that knowledge-level warrant is very frequently sufficient. Presumably, this is so either because knowledge-level warrant is quite high, or because in these very frequent cases the warrant required for action is quite low. In either case, this seems to conflict with the claim that

knowledge-level warrant is frequently necessary in these cases. For if knowledge-level warrant level is frequently necessary in these cases, this is so either because knowledge-level warrant is quite low, or because in these cases the warrant required for action is quite high. In other words, there is no reason to expect cases in which knowledge-level warrant is sufficient to be cases in which knowledge-level warrant is necessary; and there is no reason to expect cases in which knowledge-level warrant is necessary to be cases in which knowledge-level warrant is sufficient. Therefore, there is no reason to expect many cases in which knowledge-level warrant is necessary and sufficient.

Second, the stipulation that knowledge-level warrant is frequently necessary and sufficient requires treating many normal practical situations uniformly by assuming that they require exactly knowledge-level warrant (no more, no less). But, as said above, given WA, we should not expect normal practical situations to be uniform, and let alone in this way.

Consider now the second possibility, according to which, in normal cases, knowledge-level warrant is frequently close (enough) to the epistemic warrant required for acting on p:

NC** In normal cases of epistemic assessment, the degree of warrant necessary for S's knowing that p is frequently close to the warrant necessary and sufficient for the epistemic permissibility of S's acting on (the belief that) p

First, note that NC doesn't entail NC** either. For it is possible that knowledge-level warrant is very frequently sufficient and frequently necessary, but (a) the cases in which knowledge-level warrant is necessary and sufficient are not frequent, (b) the very frequent cases in which knowledge-level warrant is sufficient require much less than knowledge-level warrant, and (c) the frequent cases in which knowledge-level warrant is necessary require much more than knowledge-level warrant, so that NC is true but NC** is false.

The following table illustrates this possibility (where 'N' stands for 'knowledge-level warrant is necessary' and 'S' stands for 'knowledge-level warrant is sufficient,' '+' stands for 'much more than knowledge-level warrant' and '-' stands for 'much less than knowledge-level warrant'):

S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
N&+	N	N	N	N	N	N			
	S	S	S	S	S	S	S&-	S&-	S&-
	N&S	N&S	N&S	N&S	N&S	N&S			

Since this distribution is possible given NC, NC is still insufficient to explain the prominence of ‘know.’

Further, the arguments to the effect that NC and NC* are ad hoc apply equally well to NC**. First, NC** also relies on an unsupported frequency claim. Second, given WA we have no reason to expect normal cases to be uniform with respect to the degree of warrant required. Therefore, we should not expect most of these cases to be distributed around knowledge-level warrant.

Finally, some may think that ‘know’ is vague, and so that what ‘knowledge level’ warrant really amounts to is also vague. On this view, even if, given WA, there are many parameters which can affect the warrant, in normal cases, they do not affect the required warrant strongly enough to require more or less than what can count as knowledge-level warrant:

NC*** In normal cases of epistemic assessment, the warrant necessary and sufficient for the epistemic permissibility of S’s acting on (the belief that) p frequently falls within the borderline area for knowledge-level warrant.

However, NC*** does not follow from NC either. Suppose we grant NC but (a) the cases in which knowledge-level warrant is necessary and sufficient are not frequent, (b) the very frequent cases in which knowledge-level warrant is sufficient require much less than knowledge-level warrant, and (c) the frequent cases in which knowledge-level warrant is necessary require much more than knowledge-level warrant. NC is true but NC*** is false.

Further, the claim that the warrant required for appropriately relying on a proposition in action falls within the borderline area for knowledge level warrant is also open to the objection raised against NC, NC* and NC**. NC*** is an unsupported assumption about frequency, and given WA, we have no reason to expect NC*** to be the case. Appealing to the vagueness of ‘know’ does not help to save the proposal from ad hocness.

4. Normal and Abnormal Situations

In this section, I wish to consider whether we can defend Gerken's approach by insisting that it is only in *normal* cases of epistemic assessments that the frequency assumption applies.

An advocate of the warrant account could insist that cases in which the frequency assumption is false are (very) frequently abnormal. And even if practical parameters can combine in many different ways and affect the warrant to various degrees, it may be suggested that it is mostly in *abnormal* cases that they affect the warrant in such a way that (much) more or less than knowledge-level warrant is required.

The invocation of a distinction between normal and abnormal practical situations is problematic, though. First, the claim that 'know' is prominent in epistemic assessments of action is partially a statistical claim regarding the frequency of our actual epistemic assessments. In contrast, Normal Coincidence is a claim about a statistical fact given normal (epistemic and practical) circumstances. But if one thinks that normal practical situations are not necessarily the most frequent in the actual circumstances, then it is difficult to see how Normal Coincidence could help in explaining the fact that 'know' is (statistically) prominent in actual circumstances.

Second, on Gerken's approach, the notion of *abnormal* practical circumstances seems to be sometimes determined in terms of frequency, and sometimes not. Consider Brown's surgeon case:

SURGEON. A student is spending the day shadowing a surgeon. In the morning he observes her in clinic examining patient A who has a diseased left kidney. The decision is taken to remove it that afternoon. Later, the student observes the surgeon in theatre where patient A is lying anaesthetised on the operating table. The operation hasn't started as the surgeon is consulting the patient's notes. The student is puzzled and asks one of the nurses what's going on:

Student: I don't understand. Why is she looking at the patient's records? She was in clinic with the patient this morning. Doesn't she even know which kidney it is?

Nurse: Of course, she knows which kidney it is. But, imagine what it would be like if she removed the wrong kidney. She shouldn't operate before checking the patient's records.²⁴

²⁴ Brown, "Subject-Sensitive Invariantism," 176.

About this kind of case, Gerken writes:

For example, it should be clear enough that the practical contexts in [such cases] are abnormal due to the abnormally high stakes. Again, a broad externalist consideration motivates this assumption. It is infrequent that a single action determines life or death.²⁵

Here, it seems that the fact that a kind of action is not frequent makes it abnormal. However, Gerken continues:

Of course, surgeons face life-or-death scenarios more frequently. But this may be acknowledged by saying that surgeons are frequently in contexts with abnormally high stakes.²⁶

Here, it seems that the fact that the action is frequent does nothing to show that it is normal. As a result, it is very unclear how frequency and normality are related, and hence, it is difficult to see whether this distinction can help to show that there is no tension between the warrant account and Normal Coincidence.

Perhaps by 'normality' Gerken means the (relative) frequency for the average individual. However, it is still problematic to classify cases with high stakes (like the surgeon's case) as abnormal on the basis of the fact that they feature high stakes. This would manifest a tendency to categorize most of the practical situations in which the practical factors are such that they require a warrant (much) stronger (or weaker, if the stakes are very low) than knowledge as abnormal. But then, Normal Coincidence (or the alternative frequency assumption) would be trivially true, for given the very notion of a normal practical situation used, most situations in which (much) more or less than knowledge-level warrant is required would count as abnormal situations.

In light of these considerations, it is hard to see how an advocate of the warrant account could appeal to the distinction between normal and abnormal practical situations: either normality is not defined in terms of frequency, and hence Normal Coincidence is insufficient to explain the prominence of 'know' in actual epistemic assessments of actions; or normality is defined in terms of frequency, and then, given all the different ways in which the practical factors can combine and affect the warrant, it is quite surprising that most normal practical situations require exactly, or something very close to, knowledge-level warrant.

²⁵ Gerken, *On Folk Epistemology*, 143.

²⁶ Gerken, *On Folk Epistemology*, 144.

5. Conclusion

It is indisputable that ‘know’ (and its cognates) is prominent in ordinary epistemic assessments of action. Any satisfactory account of the epistemic norm of action must provide an explanation of this fact. Satisfying this constraint is easy for the knowledge account, but it is more challenging for alternative accounts. While rejecting the knowledge account, Gerken has developed an explanation relying on the assumption that, in normal cases of epistemic assessments, knowledge-level warrant is very frequently necessary and frequently sufficient (Normal Coincidence). On this view, conceptualising our assessments in terms of knowledge and using the word ‘know’ (and its cognates) in evaluation of action and practical reasoning would be a good heuristic. I have argued, however, that this proposal is ad hoc, for it relies on an unsupported frequency assumption which we should expect to be false given the warrant account Gerken also embraces. Further, as such, Normal Coincidence is insufficient to support the hypothesis that knowledge-based assessments of action and practical reasoning constitute a good heuristic, and alternatives to Normal Coincidence which seem better placed in this respect seem even more ad hoc.²⁷

²⁷ I would like to thank Davide Fassio, Jie Gao, Arturs Logins and Santiago Echeverri for very helpful comments on earlier versions of this paper. A previous version was presented in 2018 at the SoPhA congress in Louvain-la-Neuve (Belgium). Thanks to the audience for their comments. The work on this paper was supported by the SNSF research project ‘Rationality and Reflection’ (grant number 178039).

NOTES ON THE CONTRIBUTORS

Umut Baysan is a Lecturer in Philosophy at the University of Oxford. Previously, he worked as a University Teacher and a Postdoctoral Researcher at the University of Glasgow, where he obtained his PhD. His main areas of research are philosophy of mind and metaphysics, with research interests in epistemology. His research has appeared in journals such as *Analysis*, *Australasian Journal of Philosophy*, *Erkenntnis*, *Logos and Episteme*, *Minds and Machines*, and *Philosophical Quarterly*. Contact: umut.baysan@outlook.com.

John Biro teaches at the University of Florida. His interests are in the philosophy of language, the philosophy of mind, epistemology, metaphysics and the history of philosophy, especially in the seventeenth and eighteenth centuries. He has published papers in these areas in journals such as *The Monist*, *Australasian Journal of Philosophy*, *Philosophy and Phenomenological Research*, *Philosophical Studies*, *History of Philosophy Quarterly*, *The Review of Metaphysics* and *Hume Studies*. He is also co-editor of two collections of papers on Spinoza and one on Frege. Contact: jbiro@ufl.edu.

Fabio Lampert is a graduate student in the Department of Logic and Philosophy of Science at the University of California, Irvine. His main interests are in philosophical logic, logic, modality, epistemology, philosophy of language, metaphysics, and philosophy of religion. His publications include “Wierenga on Theism and Counterpossibles” (*Philosophical Studies*, forthcoming), “Actuality and the A Priori” (*Philosophical Studies*, 2018), “Actuality, Tableaux, and Two-Dimensional Modal Logics” (*Erkenntnis*, 2018), among others. Contact: fdalcont@uci.edu.

Miguel López-Astorga is Professor and Researcher at the University of Talca (Chile). His most recent research is focused on the human inferential activity and, especially, on the analysis of the relationships between syntax and semantics in the mental representations building processes and the derivation of conclusions. He is also interested in the reconstruction and the review of the potential of some ancient logics, particularly, Stoic logic. Contact: milopez@utalca.cl.

P.D. Magnus is a Professor of Philosophy at the University at Albany, State University of New York. He received his PhD from the University of California, San Diego and has been at Albany since 2004. His primary research is in the philosophy of science, motivated broadly by a fallibilist but non-sceptical conception of scientific knowledge. He is the author of *Scientific Enquiry and Natural Kinds: From Planets to Mallards* (Palgrave Macmillan, 2012) as well as dozens of articles. Contact: pmagnus@fecundity.com.

Timothy Perrine is a Post-Doctoral Research in the School of Philosophy at Wuhan University. He works primarily in epistemology, axiology, and philosophy of religion. He has published a number of papers in journals such as *Synthese*, *European Journal of Philosophy*, *Acta Analytica*, *Sophia*, *Faith and Philosophy* as well as edited volumes published by Oxford University Press and Cambridge University Press. Contact: timperrine@gmail.com.

Mattias Skipper is a PhD Student in the Department of Philosophy at Aarhus University. His work focuses primarily on topics in epistemology (including formal and social epistemology). He is the co-editor of *Higher-Order Evidence: New Essays* (with Asbjørn Steglich-Petersen, Oxford University Press, forthcoming), and his most recent papers include “Reconciling Enkrasia and Higher-Order Defeat” (*Erkenntnis*, 2018) and “A Higher-Order Approach to Disagreement” (with Asbjørn Steglich-Petersen and Jens Christian Bjerring, *Episteme*, 2018). Contact: filmsr@cas.au.dk.

Jacques-Henri Vollet is a postdoctoral visiting scholar in the department of philosophy at the University of Hamburg (Germany), funded by a grant from the Swiss National Science Foundation. He did his PhD at the University of Geneva (Switzerland). His areas of specialization include epistemology and the philosophy of normativity. Contact: jacquesvollet@yahoo.fr.

LOGOS & EPISTEME: AIMS & SCOPE

Logos & Episteme is a quarterly open-access international journal of epistemology that appears at the end of March, June, September, and December. Its fundamental mission is to support philosophical research on human knowledge in all its aspects, forms, types, dimensions or practices.

For this purpose, the journal publishes articles, reviews or discussion notes focused as well on problems concerning the general theory of knowledge, as on problems specific to the philosophy, methodology and ethics of science, philosophical logic, metaphilosophy, moral epistemology, epistemology of art, epistemology of religion, social or political epistemology, epistemology of communication. Studies in the history of science and of the philosophy of knowledge, or studies in the sociology of knowledge, cognitive psychology, and cognitive science are also welcome.

The journal promotes all methods, perspectives and traditions in the philosophical analysis of knowledge, from the normative to the naturalistic and experimental, and from the Anglo-American to the Continental or Eastern.

The journal accepts for publication texts in English, French and German, which satisfy the norms of clarity and rigour in exposition and argumentation.

Logos & Episteme is published and financed by the „Gheorghe Zane” Institute for Economic and Social Research of The Romanian Academy, Iasi Branch. The publication is free of any fees or charges.

For further information, please see the *Notes to Contributors*.

Contact: logosandepisteme@yahoo.com.

NOTES TO CONTRIBUTORS

1. Accepted Submissions

The journal accepts for publication articles, discussion notes and book reviews.

Please submit your manuscripts electronically at: logosandepisteme@yahoo.com. Authors will receive an e-mail confirming the submission. All subsequent correspondence with the authors will be carried via e-mail. When a paper is co-written, only one author should be identified as the corresponding author.

There are no submission fees or page charges for our journal.

2. Publication Ethics

The journal accepts for publication papers submitted exclusively to *Logos & Episteme* and not published, in whole or substantial part, elsewhere. The submitted papers should be the author's own work. All (and only) persons who have a reasonable claim to authorship must be named as co-authors.

The papers suspected of plagiarism, self-plagiarism, redundant publications, unwarranted ('honorary') authorship, unwarranted citations, omitting relevant citations, citing sources that were not read, participation in citation groups (and/or other forms of scholarly misconduct) or the papers containing racist and sexist (or any other kind of offensive, abusive, defamatory, obscene or fraudulent) opinions will be rejected. The authors will be informed about the reasons of the rejection. The editors of *Logos & Episteme* reserve the right to take any other legitimate sanctions against the authors proven of scholarly misconduct (such as refusing all future submissions belonging to these authors).

3. Paper Size

The articles should normally not exceed 12000 words in length, including footnotes and references. Articles exceeding 12000 words will be accepted only occasionally and upon a reasonable justification from their authors. The discussion notes must be no longer than 3000 words and the book reviews must not exceed 4000 words, including footnotes and references. The editors reserve the right to ask the authors to shorten their texts when necessary.

4. Manuscript Format

Manuscripts should be formatted in Rich Text Format file (*.rtf) or Microsoft Word document (*.docx) and must be double-spaced, including quotes and footnotes, in 12 point Times New Roman font. Where manuscripts contain special symbols, characters and diagrams, the authors are advised to also submit their paper in PDF format. Each page must be numbered and footnotes should be numbered consecutively in the main body of the text and appear at footer of page. For all references authors must use the Humanities style, as it is presented in *The Chicago Manual of Style*, 15th edition. Large quotations should be set off clearly, by indenting the left margin of the manuscript or by using a smaller font size. Double quotation marks should be used for direct quotations and single quotation marks should be used for quotations within quotations and for words or phrases used in a special sense.

5. Official Languages

The official languages of the journal are: English, French and German. Authors who submit papers not written in their native language are advised to have the article checked for style and grammar by a native speaker. Articles which are not linguistically acceptable may be rejected.

6. Abstract

All submitted articles must have a short abstract not exceeding 200 words in English and 3 to 6 keywords. The abstract must not contain any undefined abbreviations or unspecified references. Authors are asked to compile their manuscripts in the following order: title; abstract; keywords; main text; appendices (as appropriate); references.

7. Author's CV

A short CV including the author's affiliation and professional postal and email address must be sent in a separate file. All special acknowledgements on behalf of the authors must not appear in the submitted text and should be sent in the separate file. When the manuscript is accepted for publication in the journal, the special acknowledgement will be included in a footnote on the first page of the paper.

8. Review Process

The reason for these requests is that all articles which pass the editorial review, with the exception of articles from the invited contributors, will be subject to a strict double anonymous-review process. Therefore the authors should avoid in their manuscripts any mention to their previous work or use an impersonal or neutral form when referring to it.

The submissions will be sent to at least two reviewers recognized as specialists in their topics. The editors will take the necessary measures to assure that no conflict of interest is involved in the review process.

The review process is intended to be as quick as possible and to take no more than three months. Authors not receiving any answer during the mentioned period are kindly asked to get in contact with the editors.

The authors will be notified by the editors via e-mail about the acceptance or rejection of their papers.

The editors reserve their right to ask the authors to revise their papers and the right to require reformatting of accepted manuscripts if they do not meet the norms of the journal.

9. Acceptance of the Papers

The editorial committee has the final decision on the acceptance of the papers. Papers accepted will be published, as far as possible, in the order in which they are received and they will appear in the journal in the alphabetical order of their authors.

10. Responsibilities

Authors bear full responsibility for the contents of their own contributions. The opinions expressed in the texts published do not necessarily express the views of the editors. It is the responsibility of the author to obtain written permission for quotations from unpublished material, or for all quotations that exceed the limits provided in the copyright regulations.

11. Checking Proofs

Authors should retain a copy of their paper against which to check proofs. The final proofs will be sent to the corresponding author in PDF format. The author

must send an answer within 3 days. Only minor corrections are accepted and should be sent in a separate file as an e-mail attachment.

12. Reviews

Authors who wish to have their books reviewed in the journal should send them at the following address: Institutul de Cercetări Economice și Sociale „Gh. Zane” Academia Română, Filiala Iași, Str. Teodor Codrescu, Nr. 2, 700481, Iași, România. The authors of the books are asked to give a valid e-mail address where they will be notified concerning the publishing of a review of their book in our journal. The editors do not guarantee that all the books sent will be reviewed in the journal. The books sent for reviews will not be returned.

13. Property & Royalties

Articles accepted for publication will become the property of *Logos & Episteme* and may not be reprinted or translated without the previous notification to the editors. No manuscripts will be returned to their authors. The journal does not pay royalties.

14. Permissions

Authors have the right to use their papers in whole and in part for non-commercial purposes. They do not need to ask permission to re-publish their papers but they are kindly asked to inform the Editorial Board of their intention and to provide acknowledgement of the original publication in *Logos & Episteme*, including the title of the article, the journal name, volume, issue number, page number and year of publication. All articles are free for anybody to read and download. They can also be distributed, copied and transmitted on the web, but only for non-commercial purposes, and provided that the journal copyright is acknowledged.

15. Electronic Archives

The journal is archived on the Romanian Academy, Iasi Branch web page. The electronic archives of *Logos & Episteme* are also freely available on Philosophy Documentation Center web page.