

volume VIII ♦ issue 3
2017

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.proiectsbc.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

Michael J. SHAFFER, Internalism, Evidentialism and Appeals to Expert Knowledge.....	291
Marc Andree WEBER, Epistemic Peerhood, Likelihood, and Equal Weight..	307

DISCUSSION NOTES/ DEBATE

Philip ATKINS, Getting Gettier Right: Reply to Mizrahi.....	347
Landon D. C. ELKIND, A New Metaphysics of Sense-Data.....	359
Mark SCHROEDER, The Epistemic Consequences of Forced Choice.....	365
Notes on the Contributors.....	375
<i>Logos and Episteme</i> . Aims and Scope.....	377
Notes to Contributors.....	379

RESEARCH ARTICLES

INTERNALISM, EVIDENTIALISM AND APPEALS TO EXPERT KNOWLEDGE

Michael J. SHAFFER

ABSTRACT: Given the sheer vastness of the totality of contemporary human knowledge and our individual epistemic finitude it is commonplace for those of us who lack knowledge with respect to some proposition(s) to appeal to experts (those who do have knowledge with respect to that proposition(s)) as an epistemic resource. Of course, much ink has been spilled on this issue and so concern here will be very narrowly focused on testimony in the context of epistemological views that incorporate evidentialism and internalism, and which are either reductivist or non-reductivist in nature. Also, as the main question about testimony addressed here is whether or not testimony can provide any basic justification at all, attention will be narrowly focused on the simple case where one is presented with testimony that something is the case from only one source and on one occasion. It turns out that there are some seriously odd epistemic features of such appeals to expertise that arise both for those who intend to accept internalism, evidentialism and reductivism about justification by testimony and for those who intend to accept internalism, evidentialism and non-reductivism about justification by testimony.

KEYWORDS: testimony, expertise, internalism, evidentialism

1. Introduction

Given the sheer vastness of the totality of contemporary human knowledge and our individual epistemic finitude it is commonplace for those of us who lack knowledge with respect to some proposition(s) to appeal to experts (those who do have knowledge with respect to that proposition(s)) as an epistemic resource. Of course, much ink has been spilled on this issue and so concern here will be very narrowly focused on testimony in the context of epistemological views that incorporate evidentialism and internalism, and which are either reductivist or non-reductivist in nature.¹ Also, as the main question about testimony addressed

¹ See, for example C. A. J. Coady, *Testimony* (Oxford: Clarendon Press, 1992), Elizabeth Fricker, "The Epistemology of Testimony," *Proceeding of the Aristotelian Society* 61 (1987): 57-84, Jonathan E. Adler, "Transmitting Knowledge," *Nous* 30 (1996): 99-111, Tyler Burge, "Content Preservation," *The Philosophical Review* 102 (1993): 457-488, "Interlocution, Perception and Memory," *Philosophical Studies* 86 (1997): 21, and John Hardwig, "The Role of Trust in Knowledge," *The Journal of Philosophy* 88 (1991): 693-708.

here is whether or not testimony can provide any basic justification at all, attention will be narrowly focused here on the simple case where one is presented with testimony that something is the case from only one source and on one occasion. It turns out that there are some seriously odd epistemic features of such appeals to expertise that arise both for those who intend to accept internalism, evidentialism and reductivism about justification by testimony and for those who intend to accept internalism, evidentialism and non-reductivism about justification by testimony.

Following Conee and Feldman's insightful analysis, internalists are, typically, accessibilists and mentalists.² Briefly, internalism is the view that one's justificatory status is a function of states internal to the epistemic agent. Accessibilism is the view that epistemic agents have some sort of privileged access to those states that justify the agent's beliefs, and mentalism is the view that justifiers are mental items. Evidentialism is the view, derived from Locke, Hume and Clifford, that one should never believe anything on the basis of insufficient evidence.³ The conjunction of evidentialism and internalism then yields the view that one should never believe anything without sufficient internally accessible, mental, evidence. Non-reductivism, as it will be understood here, is just the view, derived from Reid, that testimony is a basic source of justification in the sense that it can generate justification and that such justification does not depend on knowledge of the frequency of veracity of testimony. The later condition is crucial for if testimony did require such knowledge, then it would be dependent on induction and thus would ipso facto not be a basic, justification-generating, source. Reductivism will then be understood here to be the view that testimonial justification requires knowledge of the frequency of the veracity of testimony and so on this view the justificatory status of testimony is parasitic on the justificatory status of induction. In light of the problem that these views face concerning the probativity of simple testimony it will be suggested either that externalism allows for a much more reasonable account of the epistemic role of testimony and appeals to expertise in the generation and maintenance of knowledge, or that testimony may simply not be justification-generating at all.

² Earl Conee and Richard Feldman, "Internalism Defended," in *Epistemology: Internalism and Externalism*, ed. Hilary Kornblith (Malden: Blackwell, 2001).

³ See Jonathan E. Adler, "The Ethics of Belief: Off the Wrong Track," *Midwest Studies in Philosophy* 23 (1999): 267-285, Earl Conee and Richard Feldman, *Evidentialism* (Oxford: Clarendon Press, 2004) and Trent Dougherty, ed., *Evidentialism and its Discontents* (Oxford: Oxford University Press, 2011) for perspectives on evidentialism.

2. Grave Decisions, Ignorance and Testimony

Consider a typical and simple kind of case where agent A lacks knowledge of some proposition p , or $\neg K_A p$. For example, let us suppose that John does not know if *Amanita phalloides* is poisonous or not, and that he desires to resolve this issue because he needs to know if he can safely ingest a large example of that fungi. So, relative to his question concerning the toxicity of that variety of mushroom, John wants to bring it about that either $K_{\text{John}}(\textit{Amanita phalloides}$ is poisonous) or $K_{\text{John}}\neg(\textit{Amanita phalloides}$ is poisonous). Further suppose that John is smart enough not to simply eat the mushroom in order to acquire direct evidence concerning its toxicity. As a result of his ignorance, let us then suppose that John consults a person supposed to be expert mycologist, Mike.⁴ Mike, being an expert knows that *Amanita phalloides*, the death cap, causes cyclopeptide poisoning which can result in death and which is characterized by the following gruesome pathology:

- (i) A long latent period of up to 24 hours between the ingestion of the mushrooms prior to the onset of the first symptoms.
- (ii) The occurrence of diarrhea, abdominal cramps, nausea and vomiting.
- (iii) A 24 hour period of remission of the symptoms noted in (ii), followed by
- (iv) possible liver and kidney failure, and consequent death.⁵

So, ex hypothesi, $K_{\text{Mike}}(\textit{Amanita phalloides}$ is poisonous) and $\neg K_{\text{John}}(\textit{Amanita phalloides}$ is poisonous).

Consider, however, John's epistemic position in this typical kind of appeal to expertise and where we keep in mind that John is utterly ignorant of the answer to his question. If he is, in fact, utterly ignorant of the correct answer concerning the toxicity of *Amanita phalloides*, then he presumably seeks the advice of Mike because John believes that Mike knows the correct answer, i.e. that

⁴ So in this case we have what Alvin Goldman refers to as a case of novice/expert testimony. Discussion here shall be, for the most part, limited to these sorts of cases. Also, as the concern here is with the simple question of whether single case testimony can ever justify belief, we will not be concerned with cases where the novice is faced with multiple sources that assert p . This issue is addressed at length in Alvin I. Goldman, "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research* 63 (2001): 85-110 and in George N. Schlesinger, "Why a Twice Told Tale is More Likely to Take Hold," *Philosophical Studies* 54 (1988): 141-152 and L. Jonathan Cohen, "Twice Told Tales: A Reply to Schlesinger," *Philosophical Studies* (1991): 197-200.

⁵ See Gary Lincoff. *Toxic and Hallucinogenic Mushroom Poisoning: A Handbook for Physicians and Mushroom Hunters* (New York: Van Nostrand, 1977).

Michael J. Shaffer

$B_{\text{John}}[K_{\text{Mike}}(\textit{Amanita phalloides} \text{ is poisonous}) \vee K_{\text{Mike}}\neg(\textit{Amanita phalloides} \text{ is poisonous})]$. But, it is also clearly true in this example, following Hintikka, that $\neg K_{\text{John}}K_{\text{Mike}}(\textit{Amanita phalloides} \text{ is poisonous})$ because if he did, then he would not be ignorant of the matter of the toxicity of *Amanita phalloides*. Hintikka's principle that underwrites this is as follows:

(HP) $K_A K_{BP} \supset K_{Ap}$.⁶

For the purpose of clarity, keep in mind, also, the contrapositive of HP:

(CHP) $\neg K_{Ap} \supset \neg K_A K_{BP}$.

HP is a desirable principle to satisfy as it implies that if A knows that B knows that p, then A knows that p *and* A then knows that B is an expert concerning p. However, with respect to the case we have been considering, if John knew that Mike knew that *Amanita phalloides* is poisonous then John would know that Mike knows that *Amanita phalloides* is poisonous is true. Consequently, since on the standard analysis of knowledge we cannot know what is false John would know that *Amanita phalloides* is poisonous. However, ex hypothesi, he does not know that *Amanita phalloides* is poisonous and so he does not know that Mike knows that *Amanita phalloides* is poisonous as is made especially clear by CHP.

However, it does not seem obviously necessary that John needs to satisfy HP with respect to the proposition that *Amanita phalloides* is poisonous in order to know that Mike is an expert and, hence, to make it rational to rely on Mike's testimony. So the following rationality condition for testimony is too strong:

(T1) If agent A is rational in relying on B's assertion that p, then $K_A K_{BP}$.

All that John appears to need to know is that Mike knows *whether* or not *Amanita phalloides* is poisonous, i.e. $K_{\text{John}}[K_{\text{Mike}}(\textit{Amanita phalloides} \text{ is poisonous}) \vee K_{\text{Mike}}\neg(\textit{Amanita phalloides} \text{ is poisonous})]$. Consider then the follow condition on the rationality of relying on testimony:

(T2) If agent A is rational in relying on B's assertion that p, then $K_A(K_{BP} \vee K_{B\neg p})$.

First, we must note, however, that T2 is clearly still unreasonably strong as requiring A to *know* that B knows whether or not p is true would rule out virtually every actual appeal to expert testimony on the grounds that we are rarely justified in believing such things to the degree that they count as bona fide knowledge.

⁶ Jaakko Hintikka, *Knowledge and Belief* (Ithaca: Cornell University Press, 1962), 61. Also, see Adler, "The Ethics of Belief" for a defense of a slightly modified version of HP.

So let us consider the more important and much weaker claim: $B_{\text{John}}[K_{\text{Mike}}(\textit{Amanita phalloides} \text{ is poisonous}) \vee K_{\text{Mike}}\neg(\textit{Amanita phalloides} \text{ is poisonous})]$, that John merely *believes* that Mike knows whether or not *Amanita phalloides* is poisonous. This is a reasonable rendering of the claim that John believes that Mike is an expert concerning the toxicity of *Amanita phalloides*. We can then weaken T2 as follows:

(T3) If agent A is rational in relying on B's assertion that p, then $B_A(K_{\text{BP}} \vee K_{\text{B}\neg\text{p}})$.

Suppose then that Mike were to tell John that *Amanita phalloides* is poisonous and John satisfies T3 with respect to the proposition that Mike is an expert. Recalling also that $\neg B_{\text{John}}K_{\text{Mike}}(\textit{Amanita phalloides} \text{ is poisonous})$ & $\neg B_{\text{John}}K_{\text{Mike}}\neg(\textit{Amanita phalloides} \text{ is poisonous})$, that John believes neither than Mike knows it is poisonous or that it is not poisonous, we then know, at least, the following list of relevant facts about our situation:

1. *Amanita phalloides* is poisonous.
2. $\neg K_{\text{John}}\neg(\textit{Amanita phalloides} \text{ is poisonous})$.
3. $\neg K_{\text{John}}(\textit{Amanita phalloides} \text{ is poisonous})$.
4. $K_{\text{Mike}}(\textit{Amanita phalloides} \text{ is poisonous})$.
5. $B_{\text{John}}[K_{\text{Mike}}(\textit{Amanita phalloides} \text{ is poisonous}) \vee K_{\text{Mike}}\neg(\textit{Amanita phalloides} \text{ is poisonous})]$.
6. $\neg B_{\text{John}}K_{\text{Mike}}(\textit{Amanita phalloides} \text{ is poisonous})$.
7. $\neg B_{\text{John}}K_{\text{Mike}}\neg(\textit{Amanita phalloides} \text{ is poisonous})$.
8. Mike tells John that *Amanita phalloides* is poisonous.

At least *prima facie* it should be clear that John's mere, unsupported, belief is not sufficient to rationally establish that Mike really is an expert on this matter and that T3 is thus too weak to support the contention that given 1-8, internalism and evidentialism John knows that *Amanita phalloides* is poisonous. If John is rational in appealing to Mike's expertise and thus can come to know or believe *justifiably* based on Mike's testimony, then John must have some sort of justification, either of the reductive or non-reductive sort, for his belief that Mike is really an expert on this issue.

In the immediate case at hand, it is easy to see why this is so. As far as John knows, Mike's testimony is no more liable to be correct than that of a randomly selected person from the population and, given the potential gravity of his choice, he needs to especially careful in formulating his belief. *Pace* Burge and Hardwig then, according to most reductivists it would (at least *prima facie*) seem that is not

enough that John simply *trusts* Mike without some reason to believe that Mike is knowledgeable about the toxicity or non-toxicity of *Amanita phalloides* if he is to come to have bona fide knowledge on the basis of Mike's testimony, at least if one wants to retain some standard form of evidentialism.⁷ Moreover, from a practical perspective, it would seem that John *should not* blindly trust Mike as if Mike is wrong John may die a horrible death.⁸ But, as we saw earlier, it is also unreasonable to require that John *know* that Mike is an expert. What we can reasonably expect, however, is that, in accord with evidentialism, John's belief that Mike is an expert must be justified, $J_{\text{John}}[K_{\text{Mike}}(\textit{Amanita phalloides} \text{ is poisonous}) \vee K_{\text{Mike}}\neg(\textit{Amanita phalloides} \text{ is poisonous})]$, and that his belief is so justified by sufficient evidence *e*. This yields the following rationality condition for accepting testimony:

(T4) If agent A is rational in relying on B's assertion that *p*, then there is sufficient evidence *e* available to A and relevant to $B_A(K_{\text{BP}} \vee K_{\text{B}\neg p})$ such that $J_{\text{B}}(K_{\text{BP}} \vee K_{\text{B}\neg p})$ on the basis of *e*.

But, problems then begin to creep up on those who wish to defend evidentialism, internalism and the view that testimony provides justification. These problems can be made acute by pondering the following question. What internally accessible mental items of evidence, i.e. what beliefs, could John conceivably have that would make it rational for him to regard Mike as an expert on the *specific* issue of the toxicity or non-toxicity of *Amanita phalloides*, and which would allow John to satisfy T4 with respect to that issue from his position of complete ignorance? In looking at this question, it will be shown that given internalism and evidentialism, the only plausible ways to satisfy T4 in such circumstances would require either satisfying HP, thereby rendering the appeal to expert testimony in such epistemic systems paradoxically superfluous, or by ceding the standard form of evidentialism and thus inviting incoherence. The former problem arises for reductivists because such testimony sufficiently supported by evidence will run afoul of HP. The latter problem arises for non-reductivists like Burge who hold that justificatory dependence on testimony is warranted a priori and thus cannot be evidential in the standard sense.⁹

⁷ See Burge, "Content Preservation," "Interlocution, Perception," and Hardwig, "The Role of Trust."

⁸ See Elizabeth Fricker, "Against Gullibility," in *Knowing from Words*, eds. Bimal K. Matilal and A. Chakrabarti (Dordrecht: Kluwer, 1994) for a related point.

⁹ Burge, "Content Preservation," "Interlocution, Perception."

3. The Failure of Internalist Evidentialist Reductivism

Clearly John's internal evidence cannot be that Mike has correctly answered John's question concerning the toxicity of *Amanita phalloides* in the past (either by demonstration or by having told John) and, hence, that John knows that *Amanita phalloides* is toxic in accord with HP, because then John would then already know or be justified in believing that *Amanita phalloides* is poisonous. Again, ex hypothesi, he neither knows nor is justified in believing this. Moreover, John's evidence cannot be testimonial evidence about Mike's expertise either without inviting a viscous regress with respect to Mike's testimony that *Amanita phalloides* is poisonous.¹⁰ To see the latter point suppose that Jim tells John that Mike is an expert on the matter of the toxicity of *Amanita phalloides*. So suppose that $K_{\text{John}}\{K_{\text{Jim}}[K_{\text{Mike}}(\textit{Amanita phalloides} \text{ is poisonous}) \vee K_{\text{Mike}}\neg(\textit{Amanita phalloides} \text{ is poisonous})]\}$. But on what basis can John reasonably believe what Jim says? The natural answer is that Jim possesses some relevant expertise, i.e. knowledge concerning Mike's mycological expertise and that Jim's testimony establishes this. So we might suppose that John's belief about Jim's testimony settles the issue. However, because John does not know anything about *Amanita phalloides* this prevents him from being able to justify his belief that Mike is an expert with respect to *Amanita phalloides* without further appeal to authority, his ignorance also obviously prevents him from knowing that Jim is a good judge of Mike's expertise qua *Amanita phalloides*. How, absent knowledge of the relevant matter of fact concerning *Amanita phalloides*, could John be sure that Jim knows that Mike possesses the relevant knowledge in question?

The best that we could say is that $B_{\text{John}}\{K_{\text{Jim}}[K_{\text{Mike}}(\textit{Amanita phalloides} \text{ is poisonous}) \vee K_{\text{Mike}}\neg(\textit{Amanita phalloides} \text{ is poisonous})]\}$, but this is not sufficient given the normal internalist concept of justification to make John's appeal to Mike rational without having good reasons to support John's appeal to Jim and these reasons cannot be either further testimony about Jim (say that of Diane) if one is a reductivist or reasons that imply that John knows or is justified in believing that *Amanita phalloides* is poisonous. So it seems that if John is ignorant with respect to the issue of the toxicity of *Amanita phalloides*, then it appears to be *impossible* for him to coherently appeal to testimony in a way that would yield an epistemically satisfactory, i.e. justified, answer to his query on reductivist and internalist theories of justification so understood. In other words for internalist

¹⁰ See Frederick F. Schmitt, "Justification, Sociality and Autonomy," *Synthese* 73 (1987): 43-85 for a similar point.

evidentialists of the reductivist sort it seems as if one must already know the fact in question if one is to be justified in appealing to expert authority.

To put it more clearly, what John would need to know is that Jim's (expert) testimony concerning Mike's mycological knowledge that is supposed to establish the expertise of Mike on the matter of the toxicity of *Amanita phalloides* is likely to be true. But, if John cannot tell whether Mike is really an expert because he is ignorant of the toxicity of *Amanita phalloides*, then he will also not be able to tell if Jim's claim that Mike is an expert concerning the toxicity of *Amanita phalloides* is warranted. As a result, it seems clear that testimony cannot be a basic epistemic source for internalist evidentialists of this sort as iterated appeals to testimony invite viscous regress absent some knowledge, or justified belief, concerning the proposition in question. The only apparent source of evidence that could justify John's belief that Mike is an expert on this sort of internalist evidentialist view would be for him to acquire Mike's knowledge about *Amanita phalloides* thus, paradoxically, rendering the appeal to expertise epistemically superfluous.

This crucial point can be seen most easily by employing the standard probabilistic account of confirming evidence. Suppose that John acquires confirming evidence e for the belief that Mike is an expert about $(p \vee \neg p)$. Evidence e then will have to be evidence that Mike is to be relied on in the matter of $(p \vee \neg p)$. So $JB_{\text{John}}(K_{\text{Mike}p} \vee K_{\text{Mike}\neg p})$ because, where $T_{\text{Mike}p}$ is the claim that Mike's testimony that p is true, $P(T_{\text{Mike}p} | e) > P(T_{\text{Mike}p})$ and so e confirms $T_{\text{Mike}p}$. Suppose also that Mike's testimony is evidential, such that his telling John that p is evidence for John's belief that p . So $P(p | T_{\text{Mike}p}) > P(p)$ and so $T_{\text{Mike}p}$ confirms p . If both of these things are true, then it is a trivial result of the probability calculus that $P(p | e) > P(p)$ and that e confirms p . As a result, if $JB_{\text{John}}(K_{\text{Mike}p} \vee K_{\text{Mike}\neg p})$ on the basis of evidence e , given the only legitimate candidates for what e could be HP will be satisfied in a way that makes Mike's testimony superfluous. This is not, of course, terribly surprising given a reductivist view of testimony and given Bayes' theorem.

Bayes theorem can be usefully formulated as follows: $P(p | e) = P(p)P(e | p)/P(e)$. Consider the case at hand. We have $P(p | T_{\text{Mike}p}) = P(p)P(T_{\text{Mike}p} | p)/P(T_{\text{Mike}p})$. Assuming that 8 is true $P(T_{\text{Mike}p}) = 1$, and assuming that 2 and 3 are true $P(p) = .5$. Substituting we get the following expression: $P(p | T_{\text{Mike}p}) = .5P(T_{\text{Mike}p} | p)$. In accordance with the reductivist view, it is then clear that the reasonableness of John's believing that p on the basis of Mike's testimony hinges entirely on the value of $P(T_{\text{Mike}p} | p)$ and thus on John's evidence that Mike is a reliable guide to the truth that p . By the definition of conditional probability and substitution, $P(T_{\text{Mike}p} | p) = P(T_{\text{Mike}p} \& p)/.5$. Substituting we then

get $P(p | T_{\text{Mike}p}) = P(T_{\text{Mike}p} \& p)$. Given the reductivist view that testimony is not independently probative the value of $P(T_{\text{Mike}p} \& p)$ will then reduce to $P(p)$ because on this view $T_{\text{Mike}p}$ and p will be probabilistically independent. Mike's telling John that p is true has no evidential significance with respect to p . Relative to John, who is *ex hypothesi* ignorant about p , $P(p) = .5$ and any rational, i.e. probabilistically coherent, alteration in that value will be the result of John's own direct evidence e that p is true, or $P(p | e)$. On the reductivist view all of the confirmatory work then will be done by the evidence e that John has for p simply because confirmation is transitive and the effect of $T_{\text{Mike}p}$ simply falls out because it is not independently probative. Perhaps disturbingly, this throws into question relatively substantial portions of our system of supposed knowledge as we regularly rely on epistemic authority *without* acquiring the relevant knowledge possessed by the relevant experts, without acquiring the direct justifications for our beliefs in propositions vouched for by supposed experts.

One possible and even then only partial solution open to internalist evidentialists of this sort is that John's belief that Mike is an expert qua the toxicity of *Amanita phalloides* can be supported by John's having directly acquired justified beliefs, say via perception, about Mike's expertise on closely related issues, such as his knowledge of the toxicity of, for example, *Amanita bisporigera*, which is also toxic. This sort of appeal would include appealing to inductive evidence concerning Mike's expertise as if John is really ignorant of the facts concerning the toxicity of *Amanita phalloides*, then his prior evidence concerning Mike's expertise cannot be that Mike correctly answered this question in the past as that would, again, imply that John already knew or was justified in believing that *Amanita phalloides* is toxic.

So, such appeals must involve the extrapolation of direct evidence concerning Mike's expertise on issues other than that of the toxicity or non-toxicity of *Amanita phalloides*. This sort of appeal would include appealing to inductive evidence concerning Mike's expertise as if John is really ignorant of the facts concerning the toxicity or non-toxicity of *Amanita phalloides*, then his prior evidence concerning Mike's expertise cannot be that Mike correctly answered this question in the past as that would, again, imply that John already knew that *Amanita phalloides* is toxic. However, this does not work in all cases, as John may not, in point of fact, actually have direct evidence concerning Mike's expertise on the toxicity or non-toxicity of mushrooms other than *Amanita phalloides*. Moreover, it is not at all clear that Mike's expertise concerning the toxicity or non-toxicity of *Amanita bisporigera* has evidential significance with respect to his expertise on the toxicity of *Amanita phalloides*. The sort of evidential

extrapolation principle required to underwrite this view seems dubious to say the least. It would have to take the form of some sort of principle to the effect that justified belief about A's expertise with respect to p support belief about A's expertise with respect to issues relevantly similar to p. Consider the following charitable rendering of such a principle, where $\text{Sim}(x,y)$ establishes a relevant similarity relation between propositions:

(EEP) $\text{JBAKBp} \supset \text{JBAKBq}$, for all p, q such that $\text{Sim}(p, q)$.

But, such a principle is doomed to intolerably vague with respect to the similarity relation and it seems obvious that this principle simply does not always hold. For example, Mike may never have even heard of *Amanita bisporigera* and so would know nothing about its toxicity even though we may suppose that he knows that everything there is to know with respect to *Amanita phalloides'* dangerous toxicity.

A second possible, but ultimately unsatisfactory, solution apparently open to internalist evidentialists of this sort would be to appeal to justification as coherence.¹¹ This would, in effect, appear to render moot any need for a principle like EEP. What an internalist evidentialist of this sort might be inclined to say is that the problem of establishing the bona fides of appeal to expert testimony indicates is that while internalism, mentalism and evidentialism should be retained, accessibilism can be ceded. Of course, this is due at least in part to the well-known problem of our inability to effectively compute coherence.¹² But, nevertheless, such a view would prima facie appear to allow that John is justified in believing Mike's testimony that *Amanita phalloides* is poisonous provided his believing Mike's testimony yields a more coherent total belief state than that produced by his believing that *Amanita phalloides* is not poisonous. But, it is hard to see how John can be reasonably sure that accepting Mike's testimony or any testimony does, in fact, yield the more coherent belief system if John is truly ignorant of the facts concerning the issue of the toxicity of *Amanita phalloides*.

Moreover, it is not at all clear that this sort of tactic would underwrite a general principle to the effect that we should a priori accept testimony, even of the aggregate sort, as evidentially significant. It may be true that taking testimony at face value produces the most coherent belief system, but testimony may not, even generally, be true. Whether accepting testimony as a source of evidence

¹¹ Bovens and Hartmann appear to defend such a view in Luc Bovens and Stephan Hartmann, *Bayesian Epistemology* (Oxford: Oxford University Press, 2003).

¹² See Hilary Kornblith, "The Unattainability of Coherence," in *The Current State of the Coherence Theory*, ed. John W. Bender (Dordrecht: Kluwer, 1989) and Paul Thagard, *Coherence in Thought and Action* (Cambridge: MIT Press, 2000).

produces a more coherent system of beliefs is a matter of fact to be determined a posteriori if we take coherence justification to be truth-indicative. The connection between justification and truth cannot be fixed by coherence without appeal to further evidence concerning the frequency of veracity of testimony. Nevertheless, such general facts about testimony do not appear to be able to underwrite John's believing that Mike knows the answer to the specific question of the toxicity or non-toxicity of *Amanita phalloides* without implying that John already knows that *Amanita phalloides* is toxic.

What are we then to conclude about appeals to expertise and testimony? If we accept internalism and evidentialism and reductivism, it would seem to be the case that such appeals may simply be epistemically worthless when we begin from a position of complete ignorance concerning some matter of fact. As a result, it may simply turn out that from a position of total ignorance, the only way to credibly resolve such an epistemic dearth is to seek direct evidence from a sufficiently reliable source. It also suggests, in accord with more skeptical intuitions, that we may not, in point of fact, have as much knowledge as we suppose because testimony may not provide justification. The most promising option open to internalist evidentialists who also accept reductivism might then be to claim that what we possess based on authority in far greater numbers are propositions that we merely *accept* (i.e. propositions that we entertain for pragmatic reasons without epistemic justification), especially in cases where the consequences of making a mistake are not too practically dire.¹³ What, in turn, this suggests more generally is that pragmatic and contextual factors might play a useful role in demarcating testimonial *acceptance* from justified beliefs based on testimony. To achieve the latter sort of epistemic states without falling prey to viscous regress we are required to establish, by appeal to direct evidence, that the testimony comes from an expert source in order to avoid falling prey to the inability to discriminate *Ad Verecundiam* pseudo-justifications from legitimate appeals to authority. Curiously, this does appear to render such appeals epistemically superfluous, and so shows that the allegedly overlooked centrality of testimony in epistemologies that accept these three principles will be spurious.¹⁴

In order to avoid such worries what other defenders of the epistemic basicity of testimony have done is to attempt to cast cases of pragmatic

¹³ This, of course, would be no comfort to Jamesians who reject evidentialism and argue that *belief* can be rational despite one's having insufficient evidence in sufficiently grave cases. What this discussion suggests is that James is, perhaps, confusing acceptance and belief.

¹⁴ See Robert Audi, "The Place of Testimony in the Fabric of Knowledge and Justification," *American Philosophical Quarterly* 34 (1997): 405-422 on the overlooked role of testimony.

acceptance in an artificially positive light by appeal to an exceptionally weak standard of justification, a standard so weak that it is in fact no standard at all. This brings us to the second possible option open to internalist evidentialists discussed earlier. Perhaps internalism and evidentialism can be maintained if one is willing to simply reject reductionism and accept that testimony is a basic, justification-generating, epistemic source. Doing so would seem to imply that our reliance of testimony does not require our establishing the frequency of veracity of various testimonial sources and so might well avoid the problem of the superfluosity of testimony by treating testimonial warrant as an a priori matter.

4. The Failure of Internalist Evidentialist Non-Reductivism

In this vein Burge and Hardwig would have us accept that all testimony is justificatory absent some reason to believe otherwise, absent any defeaters with respect to that testimony, and the only apparent reason they seem to do so is in order to avoid having to draw the conclusion that we possess far less knowledge than we might suppose.¹⁵ This is troubling in and of itself as it rather clearly begs the question against the skeptic, but as we shall other problems arise for non-reductivist version of internalist evidentialism as well. To begin, consider Burge's infamous trust principle:

(TP) A person is a priori entitled to accept a proposition that is taken to be presented as true and that is seemingly intelligible to him, unless there are stronger reasons not to do so.¹⁶

Now surely this would amount to a rejection of evidentialism if we read "reasons" as epistemic reasons and so would be unacceptable to the many garden-variety internalists who accept mentalism and accessibilism. Nevertheless, Burge argues that this is the essence of the non-reductivist view and that testimony is basic in the sense that it does not require appeal to other sources of justification (induction in particular) in order to provide justificatory support. The core idea behind this view is that testimony is a basic source of evidence capable not only of generating and increasing justification, but also of generating knowledge independent of empirical concerns and it seems as if this will be true presumably even if we are unaware of TP.¹⁷

¹⁵ See Burge, "Content Preservation," "Interlocution, Perception," and Hardwig, "The Role of Trust."

¹⁶ Burge, "Interlocution, Perception," 45.

¹⁷ So the view Burge endorses seems as if it is a sort of deontological internalist view and so does not include accepting accessibilism. It is not clear to me whether he endorses mentalism or not.

Putting the deeply controversial issue of whether testimonial justification can generate knowledge aside, it seems to be clear that if this position is maintained, then evidentialism, as it is ordinarily understood, must be given up and, as a result, this view amounts to nothing more than capriciousness about justification by testimony. This point can be usefully seen by once again adopting the standard probabilistic theory of justification as confirming evidence. Consider the case where we suppose that $J_{\text{John}}(K_{\text{Mike}p} \vee K_{\text{Mike}\neg p})$ because of TP and that, as a result, $J_{\text{John}p}$. But in the case we have been looking at, prior to Mike's testimony that p , $T_{\text{Mike}p}$, John is ignorant of the truth about $(p \vee \neg p)$ and so relative to John $P(p) = .5$ if he is rational. By Bayes' theorem $P(p|e) = P(p)P(e|p)/P(e)$ and to violate either this theorem or the principle of conditionalization that governs probabilistic belief updating is to invite probabilistic incoherence and is thus ipso facto irrational.¹⁸

Given the explicit constraints endorsed by non-reductivist internalist like Burge, it is easy to see that it is not possible that $T_{\text{Mike}p}$ could raise the probability of p relative to John from his stipulated state of ignorance. Consider John's state after Mike's testimony in terms of Bayes' theorem: $P(p|T_{\text{Mike}p}) = P(p)P(T_{\text{Mike}p}|p)/P(T_{\text{Mike}p})$. $P(T_{\text{Mike}p}) = 1$ if we simply stipulate 8 and assume that it is true that Mike tells John that *Amanita phalloides* is toxic. We also know that if 2 and 3 are stipulated as true then, provided he is rational, relative to John $P(p) = .5$. Substituting this information in our application of Bayes' theorem then yields: $P(p|T_{\text{Mike}p}) = .5P(T_{\text{Mike}p}|p)$. As a result, the only way that John's justification for his belief that p can alter from the initial state of ignorance is due to the posterior probability in that expression, $P(T_{\text{Mike}p}|p)$. The problem is then that the posterior probability in question is a conditional probability about the frequency of truth of Mike's testimony that p conditional on the truth of p ! But, non-reductivists are *committed* to the view that testimony is a basic form of justification and that this basicity is to be understood as the ability to generate justification *without appeal to inductive frequencies about the veracity of testimony*. So the choice is clear and forced, internalists cannot be both non-reductivists and evidentialists.¹⁹ Essentially, in rejecting reductionism non-reductivists of the internalist sort open themselves to the charge that any alteration in the probability of a proposition

¹⁸ On conditionalizing and incoherence see Paul Teller, "Conditionalization and Observation," *Synthese* 26 (1973): 218-258 and Bas Van Fraassen, *Laws and Symmetry* (Oxford: Clarendon Press, 1989).

¹⁹ The other possibility is simply to reject the standard probabilistic theories of justification, but this seems to be an unreasonably high price to pay in order to maintain the basicity of testimony.

justified solely on the basis of testimony is nothing more than an arbitrary re-assignment of a prior probability $P(p)$ to some value greater than .5 and so must ipso facto be probabilistically incoherent as it does so independent of evidence and, more worrisome yet, in doing so it also runs afoul of endorsing Moorean contradictions of the following form: I believe that p , but I have insufficient evidence that p .²⁰

If all of this weren't bad enough, accepting TP would also be an exceedingly stupid epistemic policy for someone in a situation like John's to follow. John's intention is, *ex hypothesi*, to eat the mushroom if he is told that it is not toxic and to refrain from eating it if it is toxic. So John needs to be very careful and as a result needs to adopt sufficiently stringent standards of evidence with respect to Mike's expert testimony.²¹ He needs to adopt standards that exceed those required for mere acceptance of a proposition, the state of entertaining a proposition as a basis for action or reasoning, and sufficient for at least well-justified belief.²² Bare acceptance appears to require only that one adopt a proposition as a basis for acting or reasoning, whereas rational acceptance may require only weak pragmatic justification, but rational belief requires epistemic justification, especially when there are pragmatic reasons to suppose that rational acceptance is too weak given the agent's contextual situation.²³ Bare trust then is insufficient for establishing belief for internalist evidentialists, although it may well play a role in fixing acceptance. The defenders of the epistemic basicity of testimony who are internalists appear simply beg the question against the skeptic and to concede evidentialism in order to maintain internalism and non-reductivism and they do in a way that is patently irrational from both the epistemic and pragmatic perspectives.

²⁰ See Adler, "The Ethics of Belief" on this point.

²¹ See Fricker, "Against Gullibility," on this point.

²² See L. Jonathan Cohen, *An Essay on Belief and Acceptance* (Oxford: Clarendon Press, 1992), Michael J. Shaffer, "The Privacy of Belief, Morality and Epistemic Norms," *Social Epistemology* 20 (2006): 41-54, "Three Problematic Theories of Conditional Acceptance," *Logos & Episteme* (2011): 117-125, "Doxastic Voluntarism, Epistemic Deontology and Belief-contravening Commitments," *American Philosophical Quarterly* 50 (2013): 73-82, "Epistemic Paradox and the Logic of Acceptance," *Journal of Experimental and Theoretical Artificial Intelligence* 25 (2013): 337-353, "A Thoroughly Modern Wager," *Logos & Episteme* 8 (2017): 207-231 (2017), and Robert Audi, "The A Priori Authority of Testimony," *Philosophical Issues* 14 (2004): 18-34 for discussion of the difference between belief and acceptance.

²³ Another possibility is that one might suppose that pragmatic reasons might raise S 's degree of belief that p and lower his degree of belief that $\neg p$, but Zemach has shown that it is not possible to maintain this view because practical reasons *cannot* increase the probability of a belief. See Eddy Zemach, "Pragmatic Reasons for Belief?" *Nous* 4 (1997): 525-527 for details.

5. Conclusion

Given the inadequacies of both of these views something obviously has to give. On the one hand, the most reasonable suggestion for those who wish to retain some substantive role for testimony in epistemology would perhaps be to cede internalism proper, the view that one's justificatory status is a function of states internal to the epistemic agent, and mentalism, the view that justifiers are mental items, in favor of a view that incorporates externalism, and perhaps some weaker form of evidentialism and/or accessibilism.²⁴ For example, establishing the epistemic bona fides of expert testimony is neither problematic nor is testimony incoherent or superfluous for garden-variety reliabilists. On such views Mike's testimony that *Amanita phalloides* is poisonous is reliable and John should believe it just in case he has reason to believe that Mike is reliable. He will have good accessible reasons to believe that Mike is reliable just in case he has evidence that supports the view that Mike is reliable in this regard and he will have good evidence to the effect that Mike's testimony is reliable just in case that evidence was produced by a reliable source. On the other hand, one might just be tempted to reject the view that testimony ever provides justification and that, irrespective of what Mike says, John should seek some direct evidence about the toxicity of the *Amanita* before ever considering ingesting it.

²⁴ Steup discusses the compatibility of externalism and accessibilism in Matthias Steup, "Epistemic Duty, Evidence and Internality," in *Knowledge, Truth and Duty*, ed. Matthias Steup (Oxford: Oxford University Press, 2001).

EPISTEMIC PEERHOOD, LIKELIHOOD, AND EQUAL WEIGHT

Marc Andree WEBER

ABSTRACT: Standardly, epistemic peers regarding a given matter are said to be people of equal competence who share all relevant evidence. Alternatively, one can define epistemic peers regarding a given matter as people who are equally likely to be right about that matter. I argue that a definition in terms of likelihood captures the essence of epistemic peerhood better than the standard definition or any variant of it. What is more, a likelihood definition implies the truth of the central thesis in the debate on peer disagreement, the so-called *Equal Weight View*, according to which we should give the opinions of our peers the same weight we give our own. Adopting a likelihood definition, however, does not end the debate on peer disagreement, because the alleged theoretical alternatives to the Equal Weight View, reinterpreted in the light of a likelihood definition, can in fact be shown to be compatible with this view—though the reinterpreted versions may appear less plausible than the original ones.

KEYWORDS: disagreement, peers, Adam Elga, likelihood, epistemic rationality

1. Introduction

A considerable part of this paper is based on a footnote by Adam Elga. More precisely, it is based on footnote 21 of Elga's seminal paper "Reflection and Disagreement". Since this footnote is so central to my line of argument, I quote it at full length:

My use of the term 'epistemic peer' is nonstandard. On my usage, you count your friend as an epistemic peer with respect to an about-to-be-judged claim if and only if you think that, conditional [on] the two of you disagreeing about the claim, the two of you are equally likely to be mistaken. On more standard usages, an epistemic peer is defined to be an equal with respect to such factors as 'intelligence, perspicacity, honesty, thoroughness, and other relevant epistemic virtues' (Gutting 1982, 83), 'familiarity with the evidence and arguments which bear on [the relevant] question', and 'general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias' (Kelly 2005). In defense of my use, suppose that you think that conditional on the two of you disagreeing about a claim, your friend is more likely than you to be mistaken. Then however intelligent, perspicacious, honest, thorough, well-informed, and unbiased you may think your friend is, it would seem odd to count her as an epistemic peer

with respect to that claim, at least on that occasion. You think that on the supposition that there is disagreement, she is more likely to get things wrong.¹

For my purposes, it does not matter to which passage of Elga's text the footnote is attached. What does matter is that defining epistemic peerhood in terms of *likelihood* reveals the essence of the concept. Elga's exact specification of this definition appears defective, however, and several arguments that rely upon it—among them those for his main thesis, namely that you have to take into account your peers' opinions only if their views are by and large sufficiently similar to your own—are in fact incorrect.

Besides indicating the false consequences Elga draws, however indirectly, from his peerhood definition, I explain some further implications for cases of peer disagreement, not drawn by Elga, that derive from defining peerhood in terms of likelihood. The astonishing result is that the central thesis in the debate on peer disagreement, namely the so-called *Equal Weight View*, according to which we should give the opinions of our peers the same weight we give our own, is easily seen to hold. Furthermore, this is, contrary to appearance, not the end of the debate.

In short: section 2 is a detailed reflection on the issues touched on in the footnote. Sections 3 and 5 are each dedicated to a point at which my account of peerhood departs from Elga's as presented in the footnote. The upshot is, in section 3, that his revised version of the *Equal Weight View* is unmotivated, and, in section 5, that his main thesis is wrong. In section 4, the *Equal Weight View* is proved, and it is shown that its main alternatives are in fact compatible with it. Finally, in section 6, some loose ends are picked up, and the actual limits of the applicability of the term 'epistemic peer' are pointed out.

2. Two Definitions of Epistemic Peerhood

In the footnote, Elga states, and defends, a definition of epistemic peerhood that he calls 'nonstandard'. The standard definition derives from Gutting, to whom the term 'epistemic peer' is commonly attributed, and is usually quoted in its most concise version from Kelly.² Both are mentioned by Elga. According to this standard definition, an epistemic peer is an equal with respect to a certain number

¹ Adam Elga, "Reflection and Disagreement," *Noûs* 41 (2007): 499. All brackets except the first pair are in the original.

² Gary Gutting, *Religious Belief and Religious Skepticism* (Notre Dame: University of Notre Dame Press, 1982); Thomas Kelly, "The Epistemic Significance of Disagreement," in *Oxford Studies in Epistemology: Volume 1*, eds. Tamar Gendler and John Hawthorne (Oxford: Oxford University Press, 2005).

of factors that influence a person's ability to judge a given matter. The various versions of the standard definition differ slightly as to which factors are included in the list. Following Kelly, we can roughly group the candidates into two categories, labelled, for example, *familiarity with the relevant evidence* and *general epistemic virtues*, or, as I will mostly refer to them in what follows, *well-informedness* and *competence*. Well-informedness may be taken to include, besides knowledge of all relevant facts, an awareness of all relevant lines of argument, enough time and willingness to consider them properly, and access to equipment such as scrap paper or calculators that may help to process the information. Competence may be taken to include, besides intelligence and expertise, virtues such as thoughtfulness, thoroughness, open-mindedness, intellectual courage, ingenuity, and incorruptibility. Arguably, not every potentially relevant factor can correctly be subsumed under one of the two headings. For example, it may seem implausible that well-informedness should embrace sufficient time or willingness. However, the implausibility of such a subsumption would not speak against the basic idea of defining epistemic peerhood by giving a list of factors on which candidates have to be equals.

Given this standard definition (or list definition, as I will occasionally call it), the central question in the debate on peer disagreement is whether it is reasonable to stick to one's belief when one encounters a peer who differs. It seems that the peer, being equally well-informed and competent, might just as easily be as right as oneself. If so, it is hard to justify why one should not revise one's beliefs in favour of an agnostic position when one is faced with a peer disagreement. Yet abandoning, for instance, some political or ideological belief just because some peer fails to share it appears spineless and submissive rather than deliberate and reasonable.

Before we see what Elga's non-standard definition can teach us about this puzzling situation (and what it cannot), let me clarify several notable characteristics of the standard definition, which are only implicitly mentioned by Gutting and Kelly, if at all. First, an epistemic peer with respect to one matter need not be an epistemic peer with respect to another. Our definition does not allow us to take two persons to be peers *simpliciter*; peerhood has to be *relativised to a subject (or a proposition)*. In addition, it should also be *relativised to an occasion (or a time)*, for one's degree of well-informedness concerning a given question as well as one's level of competence might change over time. One might, for example, gain extra evidence, acquire new skills, or forget formerly known

facts.³ Thus, the very same persons may be peers regarding a given proposition at one time but not at another.

Second, evidence does not include sensations and intuitions. More precisely, it does not include the *phenomenal character* of these sensory or rational seemings; it does not include *what it is like*, for example, to see the Niagara Falls, or to feel the conclusiveness of Gettier's argument. It does include everything that can be communicated, namely the content of a perception or intuition, the fact that one perceives or intuits this content, facts about one's feelings, and so on.⁴ The phenomenal character, however, is incommunicable; telling you how the Niagara Falls look does not bring about in you the same feeling you would have if you saw them. Similarly, merely emphasising how convincing Gettier's argument appears does not have the same persuasive effect as simply stating it and letting you judge for yourself. Evidence does not include qualitative experience because it is difficult to see how incommunicable experience might help to establish whether some proposition on which we disagree is true (but see the discussion on the *Extra Weight View* in section 4).

Third, 'being an equal with respect to certain factors' does not mean that a peer must be an equal with respect to *each* of the factors. It is *overall equality* that is required, not equality in every respect. This implies in particular that sameness of evidence, which is often taken to be obligatory for peerhood,⁵ is not necessarily required. Although this overall equality specification admittedly makes it more difficult to assess whether two given persons are peers, it is a natural qualification of the concept of peerhood that preserves what is valuable, namely that the

³ See Jonathan Matheson, *The Epistemic Significance of Disagreement* (New York: Palgrave Macmillan, 2015), 24.

⁴ In other words, evidence is propositional. See Timothy Williamson, *Knowledge and its Limits* (Oxford: Oxford University Press, 2000), 194–200, for a defence of this not uncontroversial view.

⁵ Cf. e.g. the role that 'full disclosure' of evidence and arguments plays in Richard Feldman, "Epistemological Puzzles about Disagreement," in *Epistemology Futures*, ed. Stephen Hetherington (Oxford: Clarendon Press, 2006), or the line of argument in Jennifer Lackey, "Disagreement and Belief Dependence: Why Numbers Matter," in *The Epistemology of Disagreement*, eds. David Christensen and Jennifer Lackey (Oxford: Oxford University Press, 2013), 254, which implicitly depends on the assumption that peerhood implies sameness of evidence. On the other hand, Feldman remarks that a peerhood definition requiring exactly identical evidential possessing would be useless because it could not be met in any real case of disagreement. See Richard Feldman, "Evidence of Evidence is Evidence," in *The Ethics of Belief*, eds. Jonathan Matheson and Rico Vitz (Oxford: Oxford University Press, 2014), 288.

relevant persons are, all things considered, in an equally good epistemic position to evaluate some proposition.

Fourth, note that peers need not be experts: two equally well-informed and competent persons might easily be two ignorant fools. While the most intriguing cases of peer disagreement are clearly those in which the peers have extensive skill and knowledge in the relevant field (for it is in these cases that confidence in our beliefs is most affected), the definition does not rule out cases in which the parties to the disagreement never had any reliable justification for their respective beliefs.

Fifth, it is sometimes argued that there are hardly any peers at all in non-idealised cases of disagreement because equality in both possessing and processing evidence is difficult to establish in real-world scenarios.⁶ However, even if this is right, it does not prove the debate on peer disagreement to be pointless. Someone fails to be my epistemic peer by being either my *epistemic superior* or my *epistemic inferior*, and if there is any difficulty in discerning which, they are at most *slightly* superior or inferior. Under the assumption that we should give our peers' opinions the same weight we give our own, it seems plausible that we should give the opinions of those who are only slightly superior only a little more weight than our own and those who are only slightly inferior only a little less. Under the assumption that we should not give our peers' opinions the same weight, however, it seems plausible that we should not give the opinions of those who are only slightly superior or inferior to us almost the same weight. Whatever the insights of the peer disagreement debate may be, they seem to carry over to other, more asymmetric cases of disagreement. Hence the debate has a clear impact even on real-world disagreement.

Finally, some authors who define peerhood in the standard way have factors on their list that I find problematic. According to Elgin, peers have to have the same background assumptions; according to Vorobej, they have to have similar and mutually intelligible manners of reasoning as well as comparably good track records.⁷ Roughly, a track record is an account of former successes and failures acquired over the course of many related performances. For example, if we often

⁶ See e.g. Nathan King, "Disagreement: What's the problem? or a Good Peer Is Hard to Find," *Philosophy and Phenomenological Research* 85 (2011).

⁷ See Catherine Elgin, "Persistent Disagreement," in *Disagreement*, eds. Richard Feldman and Ted Warfield (Oxford: Oxford University Press, 2010), 53; Mark Vorobej, "Distant Peers," *Metaphilosophy* 42 (2011), 711. More precisely, Vorobej distinguishes remote peers, who satisfy Kelly's definition but neither of his two additional constraints, from distant peers, who satisfy Kelly's definition and one of the additional constraints—it does not matter which—and perfect peers, who satisfy all conditions. He then goes on to argue that we should respond to different types of peer differently.

discuss tomorrow's weather and notice that we are almost equally reliable when it comes to guessing whether it will rain, we have a comparably good track record. The reasons why background assumptions, methodological preferences or track records should be excluded in defining peerhood closely connect to the reasons why the standard definition, in whatever variant, is problematic. To these reasons I turn now.

Recall that, according to Elga, "you count your friend as an epistemic peer (...) if and only if you think that (...) the two of you are equally likely to be mistaken."⁸ The first thing to note here is that this is not a proper definition. It does not state necessary and sufficient conditions for *being* an epistemic peer, rather it states conditions for *counting* someone as an epistemic peer. So let us instead assume that a peer is defined as *someone who is antecedently equally likely to be mistaken*, and let me note three further details concerning this definition before we proceed.

First, we need a word like 'antecedently' in the definition. Assume, for instance, that I am in a much better epistemic position regarding p than you are, but that the two of us have as yet neither formed nor exchanged any belief about p . Then my *prior* likelihood of being right about p is higher than yours. Assume further that, once we have made up our minds regarding p , we happen to agree that p is true. Then my *posterior* likelihood of being right about p is the same as yours—after all, our beliefs are identical. The term 'antecedently' signals that, in the definition, likelihood is to be understood as prior likelihood.

Second, throughout this text, 'likely to be right' is not to be understood as 'likely to hit the truth' but as 'likely to hit the view best supported by the available evidence' (the same holds *mutatis mutandis* for similar expressions). This means, for instance, that a person who arrives at a true belief due to some misinterpretation of what is in fact a deceptive body of evidence is not right; a person who correctly interprets the misleading evidence and hence arrives at a false belief, on the other hand, is right.

⁸ Although this definition is still non-standard, some authors have joined Elga in defining peerhood in terms of probability. Moffett, for example, regards a definition such as Elga's as "a very plausible account of the notion of an epistemic peer" (Marc Moffett, "Reasonable Disagreement and Rational Group Inquiry," *Episteme* 7 (2010), 357), Enoch defines a peer as "someone who is, somewhat roughly, antecedently as likely as you are to get things right (on matters of the relevant kind)" (David Enoch, "Not just a Truthometer: Taking Oneself Seriously (but Not too Seriously) in Cases of Peer Disagreement," *Mind* 119 (2010), 956), and White defines peers as equally reliable persons (Roger White, "On Treating Oneself and Others as Thermometers," *Episteme* 6 (2009), 235).

Third, if one talks about likelihood, one should specify what kind of likelihood one means. Elga refers in his definition with 'likely' to the subjective probabilities of the alleged peers. As a result, he defines under which conditions a person considers another person to be his or her peer. (In section 5, we will see where this noteworthy peculiarity leads him.) Another option would be to interpret 'likelihood' as objective probability. However, this kind of probability is ontologically obscure and epistemically difficult to access. A better alternative is to refer to the subjective probability of a neutral observer. Hence, the likelihoods are determined by the relevant credence functions that such an observer would have. Invoking a neutral observer is in fact a parallel to a conventional list definition, because by ascribing or denying peerhood to persons on the basis of such a definition, regardless of whether or not those persons take themselves to be peers, we act like impartial outsiders who aim at judging with maximal neutrality (which of course does not mean that we cannot be wrong).

A peer, I said, is someone who is antecedently equally likely to be mistaken. What reasons could one possibly have to prefer this characterisation in terms of likelihood over a list definition? Elga tells us in the footnote:

[S]uppose that you think that conditional on the two of you disagreeing about a claim, your friend is more likely than you to be mistaken. Then however intelligent, perspicacious, honest, thorough, well-informed, and unbiased you may think your friend is, it would seem odd to count her as an epistemic peer with respect to that claim, at least on that occasion.

Here, Elga lets us imagine that the two definitions come apart; that, according to his *likelihood* definition, your friend is not your peer but rather your inferior, while, according to some suitable *list* definition, your friend may very well be your peer. Then, he maintains, the result we get from his definition obviously trumps the result we get from the list definition. This seems correct, because anyone who is less likely than you to judge the truth value of some proposition correctly is *ipso facto* not your equal in judging that truth value and should thus not be regarded as your peer. In other words, equal likelihood is necessary for peerhood.

Sufficiency is harder to establish, and is not argued for by Elga (although he takes it for granted). To see a problem with the claim that an equal likelihood of being right is sufficient for peerhood, imagine two people *A* and *B* and some highly theoretical proposition *p*, which *A* considers to be true and *B* considers false. While *A* is an expert on the relevant field, *B* is merely a layman. *B*'s reason for denying *p* is her knowledge that *C* is a well-known expert regarding the matter under consideration, and that *C* believes $\sim p$. Because of this piece of testimonial

evidence, *B*'s belief is in fact antecedently as likely to be right as *A*'s. Thus we have a scenario according to which the likelihood definition tells us that *A* and *B* are peers regarding *p*, while *A* is clearly more knowledgeable than *B* and hence seems to be epistemically superior.⁹

Compare this scenario with the following variant of Christensen's well-known restaurant case:¹⁰ in a restaurant, we agree to give a 10% tip and split the bill evenly. Then you and I each calculate how much everyone has to pay. You are excellent at doing maths in your head; I, who normally perform poorly at this kind of task, use a calculator. Since we have often done computations in this fashion and compared the results, we know that it is as likely for you to make a mistake as it is for me to enter a wrong number. Are we peers?

The question, asked in this way, is ambiguous. We are clearly not peers regarding mental maths. We are, however, fully peers regarding this specific calculation. This is because whether one finds the correct result depends not only on one's calculating ability; access to useful equipment may also help. It should thus be covered by a good and detailed list definition. And since it is overall equality that matters, and not equality in every respect, a lack of competence can be compensated for by the use of technical means.

It seems *prima facie* fairly plausible not to consider access to testimonial evidence when assessing peerhood; likewise, though to a minor degree, it might appear reasonable not to take the use of a calculator into account. Concerning other resources, it seems reasonable to a still minor degree to disregard them in assessing peerhood. The crucial point now is that no categorical gap seems to lie between reliance on one kind of resource and reliance on another. Surely, sufficient time to consider the evidence properly should be on our list of factors, for insufficient time affects the respective peer statuses. What non-arbitrary justification could we have to include sufficient time but not access to sufficient scrap paper? Or access to sufficient scrap paper but not to relevant measuring equipment, or electronic means? None, it seems. If so, the apparent specialty of some resources is easily resolved. Even access to expert knowledge is, when you come to think of it, merely a means of obtaining a certain result. (And of course there is nothing special about relying on other human beings rather than

⁹ I am grateful to Stefan Reining for calling my attention to this problem. For a similar case, see also Jennifer Lackey, "A Justificationist View of Disagreement's Epistemic Significance," in *Social Epistemology*, eds. Adrian Haddock, Alan Millar, and D. Pritchard (Oxford: Oxford University Press, 2010), 302 n. 17.

¹⁰ David Christensen, "Epistemology of Disagreement: the Good News," *Philosophical Review* 116 (2007).

machines. Since computer programmes are able to answer astonishingly many questions correctly and serve increasingly often as experts, it would be ad hoc to claim a categorical difference between consulting human experts and consulting computers.) Therefore *B*, in the example above, should indeed be regarded as *A*'s peer, and 'access to expert knowledge' should be included in a good list definition.

Admittedly, this answer may appear counterintuitive. Keep in mind, however, that *B* is *A*'s peer only regarding one single proposition, *p*; with regard to closely related issues, *A* is probably much more likely to be right. The oddity of counting *A* and *B* as peers decreases once we see clearly that we ascribe or deny peerhood only with respect to an extremely narrowly limited subject matter.

(Why not define epistemic peerhood relative to a field of knowledge, rather than a proposition? The answer is that our definition would then be less significant. For if *A* and *B* disagree on *p*, and we know both that *A* is *B*'s superior on the area of knowledge to which *p* belongs, and that they are equals with respect to *p*, then the latter, more specific fact is the decisive one; it defeats the information that we get from the less specific fact. Having said this, I concede that under normal circumstances, it suffices to know whether two people are equals with respect to a certain area of knowledge, because we can quite reliably deduce from this whether or not they are peers with respect to specific propositions in that area.)

In sum: seemingly obvious counterexamples to the sufficiency of equal likelihood of being right for peerhood do not in fact show what they are intended to show; quite the contrary: they help to reveal how peerhood should be understood, and how closely it has to be tied to equal likelihood. If a list definition therefore yields a different result from the one we get from a likelihood definition, so much the worse for the list definition.

But could a thoroughly formulated list definition really yield a different result? For the sake of argument, take a list definition that includes familiarity with the relevant evidence and arguments, sufficient time and willingness to consider the evidence, access to whatever equipment is helpful in processing the evidence, intelligence, expertise, freedom from bias, sobriety, honesty, thoroughness, open-mindedness, intellectual courage, and creativity. Compare this definition to a likelihood definition. How could the two ever come apart? In order to see how, consider

DAY OF BIRTH. The day of the birth of my first child has finally arrived. My wife has been in labour for hours, and there is still no end in sight. So, on the midwife's advice, I go to a nearby restaurant for a quick meal with my in-laws. I barely eat anything. Never in my life have I been so excited. To calm my nerves, I reach for the bill and calculate what each of us would have to pay if we gave a

20% tip and split the total amount evenly. But I find it difficult to concentrate on the computation, and ultimately arrive at a sum quite different from the one my in-laws (who joined me in this game) have figured out. Recalculation proves that I have mixed up the numbers terribly.

There are several points to make here. First, something like ‘freedom from extraordinary excitement’ is not on our list of factors. Thus, if we are of the opinion that my extreme nervousness on the day of birth and the resulting computational impairment makes me my in-laws’ inferior regarding mental arithmetic, the list is incomplete. In support of regarding me as *epistemic inferior* in DAY OF BIRTH, we can adduce that, given my excitement, I am clearly less likely than my in-laws to do the calculation correctly. Hence, either DAY OF BIRTH represents a scenario in which our list definition yields a different result from a likelihood definition like Elga’s, or DAY OF BIRTH shows that our list definition is defective. And this defectiveness cannot be easily resolved. Of course, we can simply add ‘freedom from extraordinary excitement’ to our list and thereby make it immune to the specific counterexample presented by DAY OF BIRTH. But the general problem is that one can easily invent other scenarios that disclose further characteristics that are missing from our list. Philosophers are ready to come up with counterexamples that show the significance of hitherto overlooked attributes. The list is not only not exhaustive as it stands, but cannot be made so as a matter of principle, thanks to the vast variety of potentially relevant properties. A likelihood definition, on the other hand, summarises the effect a ‘perfect’ list would have. For the only plausible justification for putting further items on the list is that, by putting these items on the list, we let the respective probabilities of the supposed peers’ being right converge with each other.

A second point is that, contrary to such characteristics as ignorance of evidence, shortness of time, lack of intelligence, bias, or drunkenness, an extraordinary state of excitement influences a person’s examining abilities only very occasionally to such a degree that that person is thereby considerably less likely to get things right. Moreover, whether extreme excitement may influence a given person’s examining abilities at all is highly relative to the specific characteristics of that person. Hence including something like ‘freedom from extraordinary excitement’ on our list would make the definition too restrictive. However, a more specific description, involving, for example, the fact that a child is about to be born to a parent who tends to get extremely nervous in this kind of situation, would make our first problem more apparent, namely that we would need to add, *per impossibile*, virtually infinitely many more descriptions to our list in order to make it exhaustive.

One could also come up with the idea to include a more general characteristic than freedom from extraordinary excitement on the list, for example the ability to concentrate, and take it to exclude exactly the epistemically significant cases of extreme nervousness. The notorious difficulty with such general characteristics, however, is the vagueness of their entailment conditions. For instance, does the ability to concentrate include sobriety, or absence from test anxiety? It appears that the very same property that enables us to argue that somehow the right cases of extreme nervousness get excluded makes it hard to apply the resulting definition in concrete cases of peerhood assessment. The more general the characteristic is, the less helpful it proves to be.

(In addition, a specific problem regarding the ability to concentrate is that one could suffer from infrequent lacks of concentration in very exceptional situations, while generally having a high level of concentration. So in order to diagnose sufficiently severe lacks of concentration *before* the relevant deliberation processes start, one needs to know much more about the relevant persons than their general ability to concentrate. Equivalent problems will arise for similar characteristics.)

Speaking of unhelpfulness, we should indicate that, unlike a detailed list definition, a likelihood definition provides little guidance for judging whether or not someone is an epistemic peer.¹¹ We can easily compare the epistemic statuses of given subjects on the basis of the various properties named in a list definition; a far more abstract likelihood condition, by contrast, leaves us alone and unaided with that task. In order to assess someone's likelihood of being right concerning a specific question, we actually need to draw on their familiarity with relevant evidence and arguments, their intelligence, lack of bias, sobriety, and so on. We need the information contained in a detailed variant of the standard definition. Such a variant, however, is almost certain to be wrong, as is shown by various counterexamples, such as DAY OF BIRTH, which are ready at hand. A variant of the standard definition that uses umbrella terms such as 'well-informedness' or 'competence', on the other hand, is both potentially imprecise—there might also be cases in which it differs from a likelihood definition—and comparatively uninformative—it is, for instance, unclear whether sobriety or, for that matter, the ability to concentrate should play a role in deciding whether someone is a peer. It seems that in defining epistemic peerhood, the cost of precision is unformativeness.

It is important to note, however, that this does not pose a serious problem. A likelihood definition is as precise as could be wished and captures the entire

¹¹ See e.g. Axel Gelfert, "Who Is an Epistemic Peer?," *Logos & Episteme* 2 (2011), 512.

intent of our concept of epistemic peerhood. This is all we need in contexts in which a precise understanding of this concept matters. When it comes to applying the concept, we of course need a detailed list of properties that we can check, but no such list can itself be suitable for a definition. Moreover, in compiling such a list, we have to allow ourselves to be led by the likelihood criterion.

An analogy may be helpful. One can define *water* as the colourless stuff that falls from the skies, fills our lakes and rivers, and flows from our taps. Or one can define it as the substance whose molecular structure is H_2O . While the first definition is much more appropriate for identifying water in everyday contexts, the second is usually taken to be the correct one, the one to which we are to refer in hard cases. Similarly, the likelihood definition reveals to us the essence of peerhood, and is to be preferred in cases of doubt, while list definitions generally allow us to assess peer statuses more directly.¹²

The fact that likelihood considerations should guide us, directly or indirectly, in assessing peerhood is, by the way, the reason why background assumptions, methodological preferences and track records should preferably not be included in a list definition: they do not necessarily help to identify exactly those people who are equally likely to be right. This seems to be easy to see in the case of background assumptions and methodological preferences. After all, if you disagree with an equally well-informed and competent person on some given proposition, and the two of you track down your disagreement to a clash of deeply held assumptions, methodological or whatever, it is hard to see why your assumptions are more likely to set you on the right path than the other person's assumptions. (In fact, the matter is a bit more complex, and related to the difference between *being an epistemic peer* and *counting someone as an epistemic peer*, which I will address in section 5.)

In the case of track records, the claim that including them in the definition does not necessarily improve it may appear more surprising. To see its truth, note first that *having an equally good track record* would itself make a plausible *definiens* in a definition of epistemic peerhood. This is no coincidence: probability is often interpreted as the limit of a series of relative frequencies, and track records

¹² The analogy is suggestive but not perfect. For one thing, one could plausibly define *water* not by its actual chemical structure but by its functional role. Additionally, and relatedly, whereas the results of the rivers-and-lakes definition of water are at most slightly different from those of the H_2O definition *in the actual world*, the results of the two definitions differ considerably *in many other possible worlds*, which arguably yields some awkward consequences for the H_2O definition. Both points have no parallels in the case of the peerhood definitions, which makes the case for the likelihood definition on closer consideration far more compelling than the case for the H_2O definition in fact is.

contain such a series. A sufficiently long track record would thus give us the same information as the likelihood criterion. The track record definition, however, faces several problems. The most obvious one has its roots in the fact that the resulting probability is based only on the examination of former failures and successes, whereas in the likelihood definition we are allowed to take into account everything that might be relevant. Thus a track record definition lacks the resources to analyse the various cases of disagreement—among them those cases of enduring disagreement in fields such as philosophy or politics that interest us most—in which it is virtually impossible to get a reliable track record, not to mention a long one. To be sure, one could lower the demand for reliability and view people whose opinions on philosophical or political issues are generally reasonable as having a reliable track record, regardless of whether their beliefs are in fact true. But then surveying a track record would not amount to more than a superficial check of competence and well-informedness.

Even more importantly, the track record definition focuses on the likely effects of essential characteristics such as well-informedness and competence, not on these characteristics themselves. For this reason, it may occasionally produce wrong results: a comparably short track record might be misleading by sheer bad luck; or we might fail to see that former scenarios are not sufficiently similar to the one under consideration. For example, my track record in mental calculation, or even in doing mental calculation in states of excitement, is of no help in evaluating whether I am my in-laws' peer in DAY OF BIRTH because, in this scenario, the likelihood of my being right depends on other factors than my former performances. Moreover, as I explained above, we should take epistemic peerhood to be relative to time; this, too, does not fit well with the idea of deriving the relevant probability from the past. For example, if I begin to work out square roots in my head, my ability to do so correctly might improve rapidly, the result being that, at a certain point in time, my track record up to that point would be considerably worse than my competence. In sum, adding a track record criterion is misleading insofar as it suggests putting too much weight on past performance and too little on the specific conditions of the case under consideration.¹³

¹³ Lam defines epistemic peerhood as equal reliability, which he measures by comparing degrees of credence regarding relevant propositions to the truth values of those propositions (Barry Lam, "On the Rationality of Belief-Invariance in Light of Peer Disagreement," *Philosophical Review* 120 (2011)). Thus, Lam's definition is in fact a refined version of a track record definition, and hence faces the same difficulties. See also footnote 26 for further remarks on his conception of peerhood.

A corollary of this is that Elga's bootstrapping argument¹⁴ is flawed. Very briefly, this argument is concerned with the case of two people, you and I, say, who disagree about many not too elaborate problems that are sufficiently similar to allow for deriving reliable track records. (Elga's example is judging by eyesight which horse won one of a long series of races.) Suppose that, for each problem, it is reasonable for me to be a bit more than 50% confident that I am in fact more likely to be right than you are. Thus, as we go along, I come to regard my track record as considerably better than yours. As a consequence, I should become exceedingly confident that I am more likely to be right than you are. Yet it seems absurd to base a significantly increased certainty of being more likely to be right merely on the fact that we occasionally disagree. Therefore, so the argument goes, it cannot be reasonable for me to be a bit more than 50% confident that I am in fact more likely to be right.

Whatever other aspects of this argument might be problematic,¹⁵ it surely is not sound if we refrain from assessing other people's peer statuses primarily by their track records. For then it does not follow from the fact that my track record is markedly better than yours that I should become more confident that I am more likely to be right. Whether I am more likely to be right depends first and foremost on essential characteristics such as competence and well-informedness, whose respective degrees must remain the same throughout the whole process if the argument is supposed to make any sense. Therefore, given that in assessing someone's peer status those characteristics are more decisive than track records, my confidence of being more likely to be right should not rise, the momentous absurdity can be avoided, and the bootstrapping argument fails.

There is a third point to be made about DAY OF BIRTH, for the reason that much of what I have said so far about the likelihood definition is not quite right by Elga's lights. Look again at Elga's version of definition: "you count your friend as an epistemic peer with respect to an about-to-be-judged claim if and only if you think that, conditional [on] the two of you disagreeing about the claim, the two of you are equally likely to be mistaken." Here, Elga explicitly relativises epistemic peerhood to an "about-to-be-judged claim"—but he clearly does not relativise it to a specific time or occasion. This latter fact is surprisingly central for understanding Elga's version of the so-called *Equal Weight View*, as I will explain now.

¹⁴ Elga, "Reflection," 486–488.

¹⁵ For critical examinations of Elga's bootstrapping argument, see Duncan Pritchard, "Disagreement, Skepticism, and Track-Record Arguments," in *Disagreement and Skepticism*, ed. Diego Machuca (London: Routledge, 2013), and Jonathan Weisberg, "The Bootstrapping Problem," *Philosophy Compass* 7 (2012).

3. Elga and Equal Weight

According to a standard formulation of the *Equal Weight View* (henceforth EW), one should give the opinions of one's epistemic peers the same weight one gives one's own. This is also roughly the understanding of EW with which Elga starts on p. 484 of "Reflection and Disagreement." Four pages later, he starts refining this picture, and ends, on p. 490, with this formulation:

Equal weight view Upon finding out that an advisor disagrees, your probability that you are right should equal your prior conditional probability that you would be right. Prior to what? Prior to your thinking through the disputed issue, and finding out what the advisor thinks of it. Conditional on what? On whatever you have learned about the circumstances of the disagreement.

The reason for the refinement is the insight that the actual circumstances of the disagreement can influence your or your peers' probability of being right on a given question. To give an example, Elga lets us assume that the weather gets extremely hot, and that you know that your friend has severe problems concentrating in such circumstances. Not so you: your mental abilities are usually not affected by excessive weather conditions. If the two of you then evaluate some proposition, it may be the case that, although both of you are equally likely to be right under normal circumstances, you are more likely to be right on this particular occasion. In other words, your prior conditional probability that you will be right is not 50% but, for example, 80%. This probability is *prior* in the sense that it has to be calculated independent of what you and your friend think about the proposition that is to be evaluated; it reflects what you should¹⁶ have said about your likelihood of being right under specific conditions before you made up your mind, before the actual disagreement arose, and before you knew whether these specific conditions would actually occur. This probability is *conditional* in the sense that the specific conditions under which the disagreement actually occurs are to be taken into account.

To continue with the hot weather example, assume that you should have said prior to evaluating the claim under consideration, prior to the occurrence of a disagreement and prior to what you later learn about the specific circumstances of the disagreement that, in the event of hot weather, you should be 80% confident of being right in evaluating a proposition of a specific kind. Then Elga's version of

¹⁶ Here and at a similar location below, the Elga of "Reflection and Disagreement" would have written 'would' instead of 'should'. Enoch argues—similarly to my argument in section 5—that Elga's descriptive understanding does not capture the epistemic force we are confronted with here, and suggests a normative revision (Enoch, "Truthometer," 970–972). Enoch even reports that Elga, in conversation with him, agreed to the revision.

EW tells us that you should indeed be 80% confident when a disagreement actually arises and the weather actually gets very hot.

What has all this to do with DAY OF BIRTH? The answer is, of course, that this scenario is quite similar to the hot weather case in that my lack of concentration is limited to a special occasion: while I am normally my in-laws' equal in doing mental maths, I am their inferior on the particular day on which the birth takes place. My impairment is temporary, not permanent. As his reasons for reformulating EW suggest, Elga prefers not to deal with the possibility of such temporary impairments by relativising the definition of epistemic peerhood to specific times or occasions; he rather incorporates into EW a mechanism that prevents us from viewing our peers (in Elga's time-invariant sense) as equally likely to be right in case the circumstances of the disagreement are epistemically unfortunate either for us or for our peers.¹⁷ Hence Elga would not regard DAY OF BIRTH as a scenario that shows the advantages of the likelihood definition over the standard one; he would regard it as a scenario that indicates that the standard formulation of EW needs refining.

(Whether or not we take DAY OF BIRTH to be appropriate as a base for our arguments in favour of the likelihood definition is inessential for the dialectical force of these arguments. This is because we can easily invent an alternative scenario in which a long-lasting or even permanent change precludes me from being your peer. Assume, for example, that I have to take pills for severe depression, and that a side effect of these pills is that they reduce my attention span significantly. Since this makes it harder for me to follow lengthy lines of argument, something like 'length of attention span' should be on a standard definition's list. And again, if that is already included or entailed, other examples in which hitherto unconsidered *and enduring* characteristics play a central role are easy to find.)

The alternative to refining EW is relativising peerhood not only to an "about-to-be-judged claim" but also to a time. If we do this, we can deal with scenarios such as the hot weather case while keeping EW in its original form. And this is preferable for several reasons, which I will list in the next paragraph. Before

¹⁷ I think that this interpretation is well supported by Elga's considerations regarding the best understanding of EW. There is, however, a passage in footnote 21 that speaks against this interpretation. In defending his definition, Elga writes: "[S]uppose that you think that (...) your friend is more likely than you to be mistaken. Then however intelligent, perspicacious, honest, thorough, well-informed, and unbiased you may think your friend is, it would seem odd to count her as an epistemic peer with respect to that claim, *at least on that occasion*" (my italics). Here, Elga indeed relativises epistemic peerhood to a specific occasion. So his use of the concept seems not to be perfectly consistent in this respect.

this, I should mention that Elga gives a second argument for his version of EW, namely that it enables us to deal with peer disagreements in which we find the opinions of our supposed peers obviously irrational. The most famous example of such a disagreement is the extreme restaurant case:¹⁸ as in the normal restaurant case, you and I each calculate how much everyone has to pay. But this time the result you get is virtually impossible; it is, for example, ten times greater than the full amount stated on the bill. Do I still have to give your belief that yours is the right result the same weight as my belief that mine is correct (after all, you are my peer)? Elga observes that if we had known beforehand that we would find our peer's answer absurd, we would not have given her opinion equal weight (at least not if the case is asymmetric in that our peer does not find our answer absurd). Therefore, according to his version of EW—and his descriptive formulation: recall the 'would or should' problem addressed in footnote 16—I need not regard the probability that you are right as equal to the probability that I am right in such a case of extreme disagreement.¹⁹ In section 6, I show how we can solve this problem without revising EW.

Turning now to the reasons for preferring relativising peerhood to times over refining EW, we have to note, first, that the former proposal, in particular the formulation of EW, is considerably simpler. Second, it is somewhat arbitrary, in Elga's account, what cognitive or mental shortcomings are to be ascribed to the special circumstances of the disagreement, and what discredits someone as a peer. For example, how long-lasting must an effect of, say, pills for depression be in order to concern peerhood? What if the pills reduce my attention span for only three days (a week, a month)? There is no corresponding arbitrariness in the likelihood account, because, due to the relativisation to points in time, everything that concerns the specific circumstances of the disagreement is automatically taken into account when peerhood is ascribed or denied. Third, recall Elga's reason for favouring the likelihood definition over the standard definition: it does not primarily matter how well-informed and competent other people are concerning some proposition p ; if they are less likely to be right concerning p than we are, it appears odd to regard them as peers concerning p . The same argument can be used to show why a time-relative likelihood definition is better than a time-invariant one: it does not primarily matter whether other people are usually equally likely to be right concerning p ; if they are less likely to be right concerning p than we are on a specific occasion, it appears odd to regard them as peers concerning p on *that occasion*. Finally, our original reasons for preferring a

¹⁸ Christensen, "Good News," 199.

¹⁹ See Elga, "Reflection," 490–491.

time-relative definition still hold: one's familiarity with the relevant evidence as well as one's level of competence might change over time (and sometimes they change pretty quickly). In order to account for this, it is good to have a time-relative definition.

In this section and the previous one, I have discussed different ways of defining epistemic peerhood. Elga, in his footnote 21, states his usage of 'epistemic peer', then states the standard usage, and finally argues that his is preferable. In much more detail, I have explained the standard usage and why it is lacking (and, in many but not all respects, my arguments here are merely a specification of Elga's short remark). I have also examined Elga's definition, but have left the discussion of the discrepancy between *being an epistemic peer* and *counting someone as an epistemic peer* for section 5. This being said, the result of my discussion is that we should define epistemic peerhood and EW as follows:

Definition 1. $P_1, \dots, P_n, n \in \mathbb{N}$, are **epistemic peers** regarding a proposition p and a time t if and only if P_1, \dots, P_n are antecedently equally likely to be right when evaluating p at t .

Definition 2. The **Equal Weight View** holds that one should give the opinions of one's epistemic peers the same weight one gives one's own.

4. Why the Equal Weight View Is True—and Why This Is Not the End of the Debate

Elga takes care to emphasise that the disagreement itself does not count as evidence for whether or not the parties to the disagreement are peers. Hence we are *not* allowed to argue in the following way: "I believe that p , and I believe that you are my peer. Upon finding out that you believe $\sim p$, I have two options: either I can revise my belief that p and become agnostic about that matter; or I can revise my belief that you are my peer on the basis of your poor judgement concerning p . Which alternative is better depends on the specific proposition under consideration and the depth of my respective beliefs."

Elga claims that the second option—revising the belief that some people are our peers—is open to us only if we agree *prior* to the disclosure of a potential disagreement that we will not regard them as our peers if the disagreement actually arises. This is perfectly plausible. For assume the opposite: that I first claim that if some specific person disagrees with me, I will regard her opinion as likely to be right as I regard mine, and that I then, when the disagreement actually arises, nonchalantly downgrade her reliability. Such behaviour could hardly be

deemed internally consistent.²⁰ So either we should be prepared right from the beginning to take an occurring disagreement as a reason for not counting people as peers, and consequently deny from the outset that there could be any peers at all, or else we should not normally take the disagreement itself to be evidence for or against peerhood. Elga, for instance, allows for counting the disagreement itself as evidence only in cases in which the disagreement is rather peculiar, as it is, for example, in some variants of the extreme restaurant case. (In section 6 I will say more about taking the disagreement itself as evidence.)

According to our likelihood definition of epistemic peerhood, peers regarding some proposition p are those who are antecedently equally likely to be right about p . Observe now that any reason we might have for degrading someone's opinion is *ipso facto* a reason not to count him or her as a peer. For let us consider those reasons one by one. If a list definition were to be preferred, it would perhaps be possible to find some hidden factor that we had overlooked while compiling the list. In such a case, peers in the sense of that list definition would not be equally likely to be right about p . Yet, as I explained, we should abandon list definitions. If, secondly, the specific circumstances of the disagreement were not taken to affect the likelihoods of being right of the people involved, we would sometimes be permitted to downgrade a person's belief irrespective of his or her likelihood of being right. Yet, by relativising our definition to time, we make sure that the specific circumstances of the disagreement are taken to affect the relevant likelihoods. Thirdly, we would also be permitted to downgrade people's beliefs but not their antecedent likelihoods if the disagreement itself were relevant not only in exceptional situations like the extreme restaurant case but also quite generally. Yet, as I argued in the preceding paragraph, this condition contradicts our assumption that we consider a case of *peer* disagreement. Since the antecedents of all these conditionals have thus been ruled out, it appears that every conceivable reason for discounting someone else's opinion concerning a given question is also a reason to regard him or her as less than equally likely to be right. But if this is so, how could one *not* give the opinion of someone who is equally likely to be right as much weight as one gives one's own? Given that the peers aim at believing about p what is best supported by the evidence at hand,²¹ they are therefore compelled to give their respective beliefs

²⁰ Proponents of the Right Reasons View and the Total Evidence View, both outlined below, are likely to disagree. As my discussion of these views will show, however, they should be presented in a way that is compatible to this claim.

²¹ This is a variant of evidentialism, the view that all reasons to believe are evidential reasons, in contrast to, for example, pragmatic reasons. Evidentialism is widely acknowledged (see e.g.

the same weight. In other words, EW follows, almost trivially, from our definition of epistemic peerhood.

Does that speak against the definition? According to the several variants of the standard definition, and even according to Elga's understanding of the likelihood definition, there is room for regarding people as peers without giving their opinions equal weight. This may happen because of some characteristic of a specific scenario that is either not on our list of factors or time-relative, but which nevertheless reduces our peers' actual probability of being right. If my arguments so far are correct, however, it is arbitrary and confusing to disregard that characteristic, once it is identified, in evaluating someone's peer status. The most precise definition of peerhood is indeed not neutral with respect to the truth of EW.

The bad news (if it is bad news) is that this is not the end of the debate. It is merely a shift of focus: instead of discussing EW, we should discuss more carefully what conditions people in fact have to satisfy in order to be peers. The reason is that all well-known alternative theories to EW can be reformulated in a way that makes them compatible with EW.

Take the *Extra Weight View*. According to this view, one should give one's own opinion more weight than one gives the opinions of one's epistemic peers ('peers' understood in the sense of the standard definition).²² An extreme version of this view is the *Steadfast View*, according to which one should give the opinions of one's epistemic peers (again understood in the sense of the standard definition) no weight at all. Both the Extra Weight View and the Steadfast View come in several variants because of the various reasons one could have for

Nishi Shah, "A New Argument for Evidentialism," *The Philosophical Quarterly* 56 (2006) and Jonathan Way, "Two Arguments for Evidentialism," *The Philosophical Quarterly* 66 (2016) but not universally held (see e.g. Andrew Reisner, "The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem," *Philosophical Studies* 145 (2009)). We can, however, safely presuppose it here, because otherwise the debate on peer disagreement would not get off the ground. For assume that you are justified to believe p just because, for instance, believing p makes you significantly happier than not believing p . Obviously, this kind of justification, which is not based on evidence, is not normally undermined by learning that some peers disbelieve p , for disagreement can at most indicate that one has misevaluated the evidence, not that one has misjudged one's feelings.

²² Variants of the Extra Weight View are defended e.g. in Gideon Rosen, "Nominalism, Naturalism, Epistemic Relativism," *Noûs* 35(s15) (2001), Peter van Inwagen, "We're Right. They're Wrong," in *Disagreement*, eds. Richard Feldman and Ted Warfield (Oxford: Oxford University Press, 2010), and Michael Huemer, "Epistemological Egoism and Agent-Centered Norms," in *Evidentialism and Its Discontents*, ed. Trent Dougherty (Oxford: Oxford University Press 2011).

disregarding one's peers' opinions partially or completely. The most common variant is to adopt an agent-centred point of view and argue that *having* a certain piece of evidence is epistemically more significant than *knowing* that someone else has that piece of evidence (or a similar one). As a consequence, my intuition that p counts for more than my knowing that you have the intuition that $\sim p$.

Now, there are two ways to make sense of this. First, one could argue that it is possible that my intuition counts for more than my knowing of your contradictory intuition even if this does not make me more likely to be right. This seems to be Wedgwood's view; it comes down to denying what I took for granted in my deduction of EW, namely that we aim at believing what is most likely to be true.²³ The second way, however, is closer to the point: that my intuition counts for more just means that I am more likely to be right. In other words, even if we are equally competent and well-informed and thus are peers in terms of the standard definition (or its most elementary variant), we are not, according to this understanding, peers in terms of the likelihood definition. Therefore it does not contradict EW (and neither does its extreme variant, the Steadfast View). More generally, those adherents of the Extra Weight View who accept that we should believe what is most likely to be true do not deny that we should give the opinions of those who are equally likely to be right the same weight we give our own. They merely claim that there are, apart from characteristics such as competence and well-informedness, other factors, more closely related to one's individual perspective, that influence—and in fact increase—one's likelihood of being right considerably. In addition, they may hold that your disadvantage of not having my

²³ In a nutshell, Wedgwood's argument is as follows (see Ralph Wedgwood, *The Nature of Normativity*, (Oxford: Oxford University Press, 2009), 257–263): in order to avoid scepticism, one needs to trust at least some of the relevant intuitions that are had by reasonable people. The set of intuitions that are to be trusted without any scrutiny should not be too small—otherwise it would be insufficient as a base for non-sceptical views—but, far more importantly, it should not be too large either, because the larger the set is, the more prone to error are the theories that are built upon the intuitions it contains. Thus any *non-arbitrary* way (other ways would of course be ad hoc) of making the set smaller, up to a minimum size, is welcome. One such way, according to Wedgwood, is to trust one's own intuitions to a significantly higher degree than those of others. It is highly controversial whether this is indeed *non-arbitrary*. But even if it were, the primary motivation for the greater trust in one's own intuitions is avoidance of scepticism. Nothing in Wedgwood's line of reasoning suggests that I am presently more likely to hold a true belief about a specific proposition if I stick to my belief just because it is supported by my intuitions and undermined only by those of others. And although anti-sceptical presuppositions might be of significant epistemic value in the long run, they are not generally compatible with the aim of believing what is best supported by the available evidence and hence most likely to be true.

point of view can be compensated for by more competence and well-informedness on your side. In this case, we may be peers, in the sense of being equally likely to be right. Opponents of the Extra Weight View, on the other hand, deny that the fact that a specific assessment is one's own increases one's likelihood of being right, and hence needs (negative) compensation. Therefore, we can revise the Extra Weight View as follows:

Definition 3. The **Extra Weight View** holds that the fact that a specific assessment is one's own increases one's likelihood of being right.

This way of stating the Extra Weight View preserves what is at stake between its advocates and adversaries and adjusts it at the same time to our new conception of epistemic peerhood.²⁴

Arguably, the revised version is far less plausible than the original. For consider a disagreement between two equally competent and well-informed people: saying that each of them is justified in giving his or her own opinion more weight than another's does at least not appear contradictory; saying that each has an increased likelihood of being right, on the other hand, is hardly consistent (for instance, their likelihoods of being right cannot add up to more than 100%, but exactly this could happen if the Extra Weight View were true and if it applied symmetrically to both parties to a disagreement). This means that the only reading of definition 3 that appears defensible is as follows: Suppose *A* and *B* are equally competent and well-informed people who disagree on *p*. Then the Extra Weight View holds that, from *A*'s perspective, *A* is more likely to be right about *p* than *B*, because the fact that a specific assessment of *p* is *A*'s increases *A*'s likelihood of being right. From *B*'s perspective, on the other hand, *B* is more likely to be right about *p* than *A*, because the fact that a specific assessment of *p* is *B*'s increases *B*'s likelihood of being right. So far, so good. However, we are not so much interested in *A*'s or *B*'s perspective, but rather in the viewpoint of a neutral observer, who

²⁴ Another view that falls under my definition 3 is presented in Enoch, "Truthometer." Enoch argues that the disagreement itself has to count as evidence because of the asymmetry between *p* (a proposition that I believe prior to the disclosure of the disagreement on the basis of my original evidence) and my supposed peer's belief that $\sim p$. Whereas comparing the piece of evidence that *I believe that p* with the piece of evidence that you believe that $\sim p$ should cause me to become agnostic, given that we are peers, comparing the piece of evidence that *p* (a proposition I take to be true) with the piece of evidence that you believe that $\sim p$ need not necessarily cause me to become agnostic, according to Enoch. Although Enoch admits that the consequences of his view are precisely the same as those of an appropriate version of the Extra Weight View, he does not take his view to be a variant of the Extra Weight View. The reason is that he defines the Extra Weight View slightly more narrowly than I did.

wants to know whether A is more likely to be right about p than B or vice versa. As the answer to this cannot be ‘both’, there is no consistent neutral reading of definition 3. This is exactly what speaks against the Extra Weight View, in whatever definition; an impartial observer cannot ascribe likelihoods of being right in accordance with this view.²⁵ As this embarrassing point comes out clearer in my definition of the Extra Weight View, the task of defending this view might become noticeably harder after the shift of focus that I advocate. As we will see, similar considerations apply for the other well-known alternatives.

One of them is known as the *Right Reasons View*. According to this view, the rational thing to do in a peer disagreement is to stick to one’s opinion if one responded rightly to the original evidence and to revise one’s opinion if one responded wrongly to the original evidence.²⁶ As a consequence, it does not matter whether any disagreement, with peers or non-peers, gets disclosed after one first formed one’s belief, because what one should do depends solely on whether or not one responded correctly to the original evidence, and not on how other people evaluated this evidence. Should we therefore say that the Right Reasons View allows us to stick to our beliefs—given that we indeed hit the ones best supported by the evidence—, even if our peers favour other positions? No. For recall that our peers are those who are antecedently, i.e. before the disagreement gets disclosed, as likely to be right as we are. Then proponents of the Right Reasons View, who

²⁵ I am grateful to an anonymous reviewer for getting me to clarify this.

²⁶ Although hardly anyone defends a full-fledged version of this view, elements of it can be found in the works of several authors. Three examples: Lackey’s *Justificationist View*, according to which the prior degree of justification that one has for a particular belief is crucial for the epistemic force of disagreement about that belief, rests on an externalist notion of justification (Lackey, “Belief-Dependence,” 320). In other words, Lackey implies that one should be justified *for the right reasons*. Secondly, and even less obviously, van Wietmarschen discusses an understanding of the arguments for EW in terms of *evidential support*, a notion which is commonly regarded as being closely related to an externalist conception of justification (Han van Wietmarschen, “Peer Disagreement, Evidence, and Well-Foundedness,” *Philosophical Review* 122 (2013)). He points out that, according to such an understanding, EW fails. Thirdly, Lam, who analyses peerhood in terms of *reliability*, considers two measures of reliability, one of which (calibration) relies on a ratio of true propositions to total propositions, the other (Brier Scoring) on closeness to the truth (Lam, “Belief-Invariance”). In both cases, the reliance on truth rather than on something like (internalistically) justified response to evidence results in an externalist track record account of peerhood (cf. footnote 13). Right reasons thus play a role not in evaluating the specific disagreement under debate but in evaluating those surrounding disagreements on which the peerhood assessment is based. Given such an account, it is unsurprising that, as Lam shows, belief revision is not always called for in (non-extreme) cases of peer disagreement.

hold that, in case a disagreement occurs, one party may be right in sticking to their belief and the other wrong, should also hold that, in such a case, one party was antecedently more likely to be right than the other. To be sure, we could not have known which one, and perhaps still cannot know; but this corresponds well to the Right Reasons View, since, according to that view, we cannot know which party responded correctly to the original evidence and thus should keep hold of their opinion. In other words, since the notion of rationality that underlies the Right Reasons View is externalist, we should adopt an externalist understanding of likelihood as well. What does such an understanding look like?

Assume that I toss a coin. What is the probability of its coming up heads? Here is a surprising answer: either the world is such that the coin will come up heads. Then this event will definitely occur; hence, its probability is 1. Or the world is such that the coin will come up tails. Then *this* event will definitely occur, and the probability of the coin coming up heads is 0.²⁷ In other words, in a world in which each proposition is either true or false, the probability that a proposition is true is always either 0 or 1. A notion of probability that allows for values between 0 and 1 for future events is grounded in our limited knowledge about how things will turn out, not in the way they will actually turn out. In this sense, it is internalist. If I say “The probability that it will rain tomorrow is 30%,” I do not mean that it is not completely certain, given today’s worldwide weather conditions, whether or not it will rain tomorrow; what I mean is that in light of all the evidence I have at hand and can evaluate properly, my rational degree of credence is 30%. An externalist notion of probability cannot serve the purpose of reporting such a rational degree of credence between 0 and 1, caused by insufficient information. In the same way an externalist notion of rationality cannot account for the reasonableness of believing, on the basis of misleading evidence (namely one’s peer’s misevaluation of the original evidence), what is in fact not supported by the original evidence. This analogicity of the two externalist understandings carries over to the respective views so that we arrive at this definition:

Definition 4. The **Right Reasons View** holds that two persons are peers regarding some proposition *p* if and only if they will either both correctly evaluate the

²⁷ In what follows, I ignore theories of objective probability, according to which the world may be such that the probability of the coin coming up heads is between 0 and 1, because taking those theories into account would merely complicate matters, but not cause essential revisions. The reason is that invoking objective probabilities only has consequences for the values that externalist probabilities can take, but not for the difference between externalist and internalist probabilities. It is this difference on which my reasoning here relies.

evidence regarding p , or if they will both incorrectly evaluate the evidence regarding p , i.e. if their respective likelihoods of being right are either both 1 or both 0 (whereby likelihood is to be understood in an externalist way).

Like the original version, this reformulated one claims that the actual occurrence of a disagreement is of no epistemic significance; one should revise one's view if one is wrong or has probability 0 of being right, but not if someone else thinks differently. In sum, the Right Reasons View can be made compatible with EW by claiming that, in order to be peers, two persons have to have the same externalist probability of being right.

Relying on externalist probabilities, however, is a very strange thing to do. We are not used to using them, and there is a good reason for that. Talk of probabilities is meaningful insofar as it enables us to distinguish degrees of uncertainty. It does not matter for this aim whether we are uncertain because there is some definite fact of the matter that we do not (yet) know, or because there is no fact of the matter at all (or not yet). To use probability talk to indicate whether some fact holds true is a confusing and needlessly complicated way of speaking. Since the revised variant of the Right Reasons View forces us to interpret the notion of likelihood, as it figures in the peerhood definition, in an externalist fashion, it forces us to adopt that confusing and overly complicated way of speaking.

A third view often seen as an alternative to EW is the *Total Evidence View*, which is most prominently defended by Kelly.²⁸ According to this view, "what is reasonable to believe [in a peer disagreement scenario] depends on both the original, first-order evidence as well as on the higher-order evidence that is afforded by the fact that one's peers [understood in the sense of the standard definition] believe as they do."²⁹ The first-order evidence regarding some non-doxastic proposition p comprises all evidence regarding p except evidence regarding what others believe about p . Second-order evidence regarding p is evidence regarding what others believe about p . Third-order evidence regarding p is evidence regarding what others believe about what others believe about p . And so on. Higher-order evidence is evidence that is not first-order. The idea behind the Total Evidence View is that, although the disclosure of a disagreement gives us higher-order evidence that supports suspension of judgement, acquiring this higher-order evidence gives us no reason to completely disregard the first-order evidence on the basis of which we formed our original belief; what is reasonable

²⁸ In Thomas Kelly, "Peer Disagreement and Higher-Order Evidence," in *Disagreement*, eds. Richard Feldman and Ted Warfield (Oxford: Oxford University Press, 2010).

²⁹ Kelly, "Peer Disagreement," 142.

to believe depends on *both* sorts of evidence. As a result, it may—but need not—happen that one should not suspend judgement in the case of a peer disagreement because the first-order evidence clearly supports one position over the other.³⁰ There is thus no golden rule on what to do in the case of a disagreement with someone who is equally competent and well-informed. The Total Evidence View allows for splitting the difference if the first-order evidence is sufficiently inconclusive, but otherwise it prescribes sticking to one's belief. So it would crucially depend on the quality of the first-order evidence whether one should accept some equally competent and well-informed person as someone who is equally likely to be right in case of a disagreement. If the first-order evidence appears strong, one should rather trust one's own assessment of it than another person's opinion, and since one knows this in advance, before a disagreement arises, one should, according to the Total Evidence View, refrain from considering other people as one's peers from the outset. Hence we can reformulate the Total Evidence View as follows:

Definition 5. The **Total Evidence View** holds that whether two persons are peers regarding some proposition *p* depends, among other factors such as their competence or well-informedness, on the first-order evidence they have for *p*.

There may be cases in which the first-order evidence plays virtually no role at all in deciding who is a peer, and there may be those in which the quality of the first-order evidence makes it more or less impossible that there are any peers at all. This version of the Total Evidence View is compatible with EW.

As with the other alternatives, however, the revised version might appear harder to believe than the original. That the first-order evidence adds some information, over and above what the higher-order evidence tells us, to *what it is reasonable to believe* seems *prima facie* more plausible than that it adds some information, over and above what a person's competence, well-informedness etc. tell us, to *how likely that person is to be right*. For recall that people's likelihoods of being right on a certain matter should be determined *before* they make up their minds on that matter. However, claiming that the first-order evidence influences these likelihoods in effect means that their likelihoods depend on the positions they will take: if they take a position that is clearly more strongly supported by the first-order evidence, their likelihood will increase; if they take a position that is clearly less strongly supported, their likelihood will decrease. Hence the revised

³⁰ See Thomas Kelly, "Disagreement and the Burdens of Judgment," in *The Epistemology of Disagreement*, eds. David Christensen and Jennifer Lackey (Oxford: Oxford University Press, 2013), 50–51.

version of the Total Evidence View seems to be incompatible with our preferred understanding of the likelihood definition.

One might suspect that not only alternative theories of EW are in need of reformulation owing to our revision of the peerhood definition, but also EW itself. I seriously doubt, however, that there is any plausible way of formulating a version of EW that is not entailed by the likelihood definition without failing to capture EW's original intent. Among the implausible ways, one is the way of exclusion. According to this way, EW is defined as the thesis that the Extra Weight View, the Right Reasons View and the Total Evidence View are all wrong. This way is implausible not only because it is unlikely that those three views are the only meaningful rivals one could imagine, but also because it completely leaves in the dark what EW actually says and why we should have any interest in it. Another implausible way is that of refocusing. Here, EW states that, under normal conditions, it is not too hard to find disagreeing interlocutors who are equally likely to be right. (More often than not, one could add, it suffices if they are equally competent and well-informed.) The idea here stems from the observation that, in all three rival theories that I have discussed, peerhood depends on some special condition (point of view; actually correct response to evidence; nature of first-order evidence) that is quite unlike those factors listed in standard definitions. This arguably makes peerhood harder to come by. However, this way is in fact no better than the former exactly because it merely spells out a common strand in the three rival theories and does not add any substantial content. Moreover, an adherent of the true intent of EW should have no problems in allowing as a possibility that disagreeing peers are hard to find. Yet another implausible way of reformulating EW could be called the way of ignorance. According to this way, EW is defined as the thesis that we should give the opinions of those who are equally competent and well-informed the same weight we give our own. Here, there is simply no mention of epistemic peers in the likelihood sense. This way is implausible because our insights from the previous section are ignored and, as a result, a deficient notion of epistemic peerhood is still, though implicitly, in use. In sum, several conceivable attempts to adjust EW to our new understanding of peerhood prove to be inadequate. This is not surprising: what is essential in EW is that it connects the concept of epistemic peerhood with the prescription to give the opinions of those who fall under this concept equal weight. Since epistemic peerhood should be understood in terms of likelihood, it is virtually impossible to conceive of a plausible reformulation of EW whose truth does not follow almost immediately from our likelihood definition.

In the preceding paragraphs, I have not taken pains to present *conclusive* arguments against the several alternative theories to EW that I have discussed (nor will I do that in what follows); the most I have done is to indicate that they may appear less plausible after revision. My aim here has chiefly been to explore whether the most prominent alternatives to EW can be reformulated in a way that is *compatible* with the likelihood definition of epistemic peerhood and, *a fortiori*, with EW itself. The result is that they can indeed be appropriately reformulated; although they are arguably less plausible after revision, they cannot be ruled out as a consequence of adopting a likelihood definition. What we gain from our new framework, then, is that we see old theories in a new light, and that this helps us to understand their respective entanglements better. What is more, the point of view presented here is not merely some new one; since it is, if I am right, the one that arises from the preferable definition of epistemic peerhood, this point of view is the most suitable.

I have not taken pains to present conclusive arguments against the alternatives to EW that I have discussed; but I will, in the next section, argue—conclusively, I hope—against a marginalisation of EW that Elga, who accepts EW (or his refined version of it), puts forward in “Reflection and Disagreement.” In doing so, I shall elaborate further on what understanding of the likelihood definition is appropriate.

5. Clusters of Controversy

If what I have said so far is correct, we should not discuss whether EW is true—for it clearly is—but rather what circumstances, apart from well-known ones such as lack of competence or information, might prevent someone from being equally likely to be right. Does disagreement on some very fundamental principles, for instance, suffice for not counting someone as a peer regarding matters stemming from those principles? Elga thinks it does; more particularly, he holds that in a ‘cluster of controversy,’ that is an extended field of related issues on which disagreement prevails, one is unable to determine whether the people one disagrees with are one’s peers. To see what this means, consider

FOUR-DIMENSIONALISM. I am of the opinion that four-dimensionalism is an extremely useful and well-founded philosophical theory. Our metaphysical worldview gets so much more elegant and straightforward once we adopt it! Unsurprisingly, I find highly plausible many theses that fit well with my adherence to four-dimensionalism: that persons just are mereological sums of person stages; that identity can be contingent; that eternalism is to be preferred over presentism, counterpart theory over accounts of transworld identity, semantic approaches to vagueness over ontic or epistemic ones; and so on. You,

on the other hand, find hardly any of this plausible. You argue that we essentially conceive the world as consisting of enduring material objects (among them, most notably, ourselves), and that the structure of our thought about the world cannot be adequately captured by a metaphysical framework so alien to experience and common sense. It is thus not astonishing that although we are both (as we are happy to admit to each other) very able and well-read philosophers, the two of us favour completely different theories in many areas of modern philosophy.

FOUR-DIMENSIONALISM confronts us with a cluster of controversy about broadly metaphysical issues. If you and I are peers in this scenario, we must, according to EW, suspend judgement about whether four-dimensionalism is true. Of course, we could deny that we are peers by adopting one of the alleged rivals to EW discussed in the previous section. Elga, however, rejects these views. He argues instead that the mere fact that our disagreement is not isolated but concerns a whole bunch of related claims makes it impossible for me to count you as a peer. This is the case because, in order to judge whether you are my peer, I normally consider how reliable your beliefs about related matters are. For example, if I am about to judge whether you are my peer concerning some multiplication problem, my evidence is how good you generally are when it comes to calculating. In FOUR-DIMENSIONALISM, however, I cannot tell how reliable your beliefs about related matters are, as our controversy extends to related matters as well. (In fact, I think it goes as deep as to the question of what demands a good philosophical theory should satisfy.) To suppose that your beliefs are misled because they differ from mine would beg the question. Neither can I suppose, according to Elga, that you are as competent a metaphysician as I am, and that your beliefs in this area are, for that reason, as reliable as mine, because I simply lack the resources to judge whether you are as competent as me. How could I assess your metaphysical competence if your idea of how closely our metaphysical theories should resemble the way we actually conceive of the world differs so much from mine? Therefore, Elga denies that, in cases like FOUR-DIMENSIONALISM, the parties to the disagreement think that they are peers.³¹

(In fact, the last step of the argument is a bit too fast. What is lacking is the premise that mere ignorance about other people's peer statuses does not justify regarding them as peers; a positive doxastic attitude is needed. If one believes neither that some other people are epistemically superior nor that they are epistemically inferior, one would not thereby have a reason to count them as

³¹ See Elga, "Reflection," 492–497. Elga's own example concerns abortion, not four-dimensionalism. All further arguments in this text apply equally well to both cases.

equally likely to be right. This appears most plausible if there is no conceivable way of comparing two people's respective peer statuses. In this case, we do not just suffer from a removable lack of knowledge; we rather face a scenario in which the question of whether those people are peers becomes meaningless. Perhaps Elga should be interpreted as believing that clusters of controversy do constitute such an unbridgeable gulf for peerhood ascriptions.)³²

It should be emphasised that Elga, in presenting his line of reasoning, consistently uses expressions like "someone counts another one as a peer," "someone thinks another one is equally likely to get things right," or "someone has a certain opinion about another one's abilities." He says neither anything like "someone *is* another one's peer" nor anything like "someone *should* count another one as a peer." Throughout the whole argument, questions about whether or not certain people are peers are asked, or answered, only from the point of view of these people. There is no 'view from above,' no third-person perspective; neither is there a normative dimension involved. This is quite surprising, because, for one thing, the objective, third-person point of view appears to be considerably more relevant than the first-person one. After all, the parties to a disagreement are usually much more prone to error in judging the other parties' abilities than an impartial observer who specifically concentrates on peerhood issues. For another thing, we cannot avoid bringing in a normative dimension sooner or later because our aim in epistemology is not to describe what kinds of beliefs people actually tend to hold under specific circumstances but to explore what kinds of beliefs they should hold under such circumstances. What interests us is not under what circumstances people generally happen to count others as peers but under what circumstances they are justified in doing so.

So, in short, we can ask peerhood questions in three ways: first-person, third-person, and normative. How are these three ways connected? What are Elga's reasons for asking only the first-person questions? And is he right in doing so?

Recall Elga's definition of peerhood in the footnote: "On my usage, you count your friend as an epistemic peer with respect to an about-to-be-judged

³² Most philosophers seem to accept the premise quite generally anyway. See e.g. Enoch, "Truthometer," 956. An exception is Vulich, whose argument, however, is based on a reformulation of EW, according to which it is not suspension of judgement that is called for in the case of a peer disagreement but reconsideration (Richard Vulich, "Peer-Hood," *Logos & Episteme* 2 (2011)). A more conclusive (and, at second glance, equally pertinent) objection to the premise is King's argument to the point that lacks of clarity concerning peer statuses raise similar epistemic problems as cases of peer disagreement (King, "Disagreement," 267–269).

claim if and only if you think that, conditional [on] the two of you disagreeing about the claim, the two of you are equally likely to be mistaken.” Elga defines *epistemic peer* only for first-person, descriptive contexts.³³ His reasons for this are given in another footnote (number 14), in which he points out why he describes a peer disagreement problem in a certain way:

Note that in setting up the problem, the initial assumption is that you *count* your friend as your epistemic peer. That contrasts with some presentations, in which the initial assumption is that your friend *is* your epistemic peer. The former assumption is appropriate, however. For example, one sometimes is reasonable in thinking wise advisors to be foolish. Evidence, after all, can be misleading. In such cases, one is reasonable in being guided by one’s *assessments* of the advisor’s ability, even if those assessments are in fact incorrect.³⁴

Here, Elga draws our attention to cases in which one’s judgement that another one is one’s peer is wrong because of, for instance, misleading evidence. In these cases, he claims, we are nevertheless right in considering the other one a peer because our judgement is all we can rely on. One cannot be guided by how things are but only by how one takes things to be. For that reason, it does not matter whether peerhood actually holds; reasons for belief revision can arise only if one *thinks* that peerhood holds.³⁵

There are at least two points to make here. First, while it is certainly right that we should give equal weight to the opinions of those whom we have reason to count as our peers, regardless of whether they actually are, it is not clear why this fact presents a reason not to define epistemic peerhood in the way I did. With my definition in place, we can equally well set up peer disagreement scenarios in which we assume that the parties to the disagreement count themselves as peers. (Elga does not explicitly deny this, but refrains nevertheless from giving a full-fledged definition of peerhood.)

The second point is that Elga’s considerations do not speak against a normative account, that is against saying that we should give equal weight to the opinions of those whom we *should* count as our peers (where the rational normativity signified by *should* is internalist). In fact, Elga himself brings in a normative element when writing that “one is *reasonable* in being guided by one’s

³³ The wording in the definition is of course second-person. What matters here, however, is that the ascriber is not abstracted away, and that we thus cannot extract necessary and sufficient conditions for being a peer *simpliciter*.

³⁴ Elga, “Reflection,” 499. All italics are in the original.

³⁵ See Nicholas Tebben, “Peer Disagreement and the Limits of Coherent Error Attribution,” *Logos & Episteme* 4 (2013), 179–180, and King, “Disagreement,” 262.

assessments.” This is a natural move because beliefs obviously differ with regard to their reasonableness. And the belief that some people are one’s peers is reasonable precisely if one is justified in counting them as peers. So the idea that the crucial question, from an epistemologist’s point of view, is whether one should count someone as a peer (and not whether one would do so) might be regarded as a natural specification of what is meant here by “counting someone as a peer.”

Perhaps Elga had something like this in mind. Perhaps he thought that it goes without saying that one normally counts people as peers if and only if one should do so. Perhaps he chose his specific wording only to emphasise that it is always from a first-person point of view that peerhood is ascribed or withheld, and not to emphasise that a normative understanding is out of place. Be that as it may; the effect of his consistent avoidance of normative vocabulary and his constant adherence to first-person language as far down as to the definition of peerhood is that his argumentation to the point that people involved in a cluster of controversy cannot be peers seems more convincing than it otherwise would. The reason is that in asking the normative question, we often feel forced to reflect more deeply on the matter under consideration, and may in this indirect way eventually come to take on what is in fact a third-person perspective. To be sure, in doing so, we can never acquire the point of view of an omniscient observer; but we can, up to a certain extent, abstract from what we actually believe. In other words, we may widen the first-person perspective, which we cannot cast off in all real cases of disagreement, by reminding ourselves that we can tackle problems from different ends, and that only trying to do so might reveal us a sufficiently broad view. Let me give an example.

Concerning FOUR-DIMENSIONALISM, I have argued on Elga’s behalf that I lack the resources to judge whether you are my peer because assessing the quality of your beliefs on related matters is out of question for me for the simple reason that your beliefs on related matters are part of the same cluster of controversy. The consequence, according to Elga, is that I have no determinate opinion about whether you are my peer. This may seem a reasonable description of what one would think in this scenario. But is it a reasonable description of what one *should* think here? If the question is asked that way, other considerations may arise: am I really justified in ignoring your opinion just because our beliefs differ so extensively? How could it be that I am rationally required to hold metaphysical views so completely different from those that you are rationally required to hold (recall that the scenario is symmetric)? Is it really impossible to assess, however imperfectly, the quality of your reasoning in comparison to mine? In order to approach these questions, imagine someone who is neutral on the whole cluster of

issues surrounding four-dimensionalism. Such a person would not care about how extended our controversy is because her ignorance of the entire cluster makes it impossible from the outset for her to assess how reliable our beliefs on related matters are. In trying to assess our peer status, however, she would naturally draw on less specific pieces of evidence such as general intelligence, lucidity of reasoning, or professional reputation. (In a way, she would pay less attention to track records and more attention to the characteristics that usually cause such records. This observation fits well with the criticism of the track record criterion at the end of section 2. Recall also that my critique of the list definition is perfectly compatible with the fact that we should usually check people's peerhood statuses with the help of a detailed list of properties) If the result is that, to the best of her knowledge, the two of us are on the whole equally good metaphysicians, then she is justified in regarding us as peers and, consequently, in becoming agnostic about whether four-dimensionalism is true. We, as parties to the disagreement, should do the same: even if we have to set aside many of our deeply held metaphysical beliefs, there is enough left to enable us to assess, however roughly, whether or not we are equally competent and well-informed, and hence whether or not we are peers. It seems odd that the improbability of this assessment should be the crucial reason for me to stick to my original belief that four-dimensionalism is well-founded (and for you to stick to your original belief that it is not).

In sum, once we realise that, from the viewpoint of people who do not have to set aside many deeply held beliefs, the comprehensiveness of our disagreement need not prevent assessment of our peer status, we should join them in forming a belief about our peer status from a broader point of view. Therefore, according to my line of argument, and *pace* Elga, we cannot generally deny a peer status to people with whom we disagree extensively. A corollary of this is that a list of factors by means of which we assess epistemic peerhood must not include background assumptions or methodological preferences.

6. The Limits of Rationality

So far, I have argued that a thorough understanding of the concept of epistemic peerhood results in a likelihood definition, which in turn leads to immediate acceptance of EW, and that the mere fact that a disagreement is widespread and deep-rooted does not rule out the fact that the parties to the disagreement are peers. Yet how widespread and deep-rooted could a peer disagreement maximally be? Are there any limits? Or, to put the question differently, are there any non-

trivial circumstances that generally allow us not to regard equally well-informed and competent people as peers?

The lesson we can learn from Elga's line of reasoning is that we need some common ground in order to compare people with respect to their probability of being right on a certain matter; once we set aside too much, we are at a loss concerning whether they are, for example, equally competent. Contrary to what Elga thinks, clusters of controversy are not sufficient for rendering us so uncertain, for we can still rely upon comparatively unspecific but sufficiently informative characteristics such as intelligence, thoroughness, or freedom from bias. As we have seen, we are in fact even required to rely upon these factors because, first, what interests us is not whether we would count someone as equally likely to be right but whether we should do so, and, second, this obligation cannot be fulfilled by ignoring evidence such as reliable information about those characteristics. We would encounter serious problems, however, if we were not able to acquire such reliable information. Then we would indeed have no idea whether we should count someone as our peer. To see how this can happen, consider

THE AWKWARD COLLEAGUE. I remember that there is a meeting tomorrow. Unfortunately, I have forgotten who will take part. So I ask a colleague of mine, Jane, whom I know to be a reliable and in fact pretty smart person. She says that, apart from me and her, only Simon and Sue will be there, so in total we will be five people. "Wait a moment," I say, "you and I are two people, and Simon and Sue are also two people; that makes four, so four is the number of people who will be there." Jane shakes her head in disbelief and asserts, "No, two and two equals five."

The situation here is supposed to be symmetric: Jane is as baffled as I am when she notices that we differ about what two and two adds up to. So I cannot just refer to the profoundness of my belief that $2+2=4$, or the strength of my justification for it, in order to discredit her as a peer.³⁶ But is there a way of establishing the opposite, namely that she is indeed as likely as I am to be right? (A consequence thereof would be that I have to suspend judgement on whether $2+2=4$.)

In assessing whether Jane is my peer in elementary maths, I cannot draw on how reliable her beliefs about other simple computations are. Even if the results she arrives at are frequently different from mine, I am not allowed, according to the argument in section 5, to conclude that her beliefs are erroneous solely from

³⁶ Taking symmetry seriously also rules out Christensen's idea that common-sense checking resolves the problem (Christensen, "Good News," 199–201). This point is also made by Tomas Bogardus, "A Vindication of the Equal-Weight View," *Episteme* 6 (2009), 329.

the fact that we are tangled in a cluster of controversy. If, on the other hand, our disagreement is restricted to very few cases (or perhaps even to the one concerning $2+2$), this is not by itself a proof of the reliability of her beliefs because track records—and nothing other do I compile when I compare her results with those that I regard as true—are, as we have seen, only probable but not indefeasible effects of what is actually essential, namely such characteristics as competence and well-informedness. Yet it is exactly her competence that I have reason to doubt when faced with her apparently insane belief that $2+2=5$. This single belief is, given all I know about addition, so irrational that it undermines any argument that deduces from her being almost always correct that she is competent in adding numbers. Consider, for instance, the question of whether it is rational for someone who believes that $2+2=5$ also to believe that $3+3=6$. I do not think that our understanding of rationality allows us either an affirmative or a negative answer to this question; we simply have no idea, and cannot have any idea, what way of adding numbers would be rational under the presumption that $2+2=5$. Therefore it does not matter, under this presumption, how good our track record is; any track record is made insignificant by a disagreement on such a basic level.

These considerations suggest that, in extreme cases of disagreement, the disagreement itself can indeed count as evidence for whether or not the parties to the disagreement are peers. Elga would agree, as I mentioned at the beginning of section 4. His argument is that if we agree prior to the disclosure of the disagreement that we will not count someone as a peer in case we find her opinion utterly insane, then actually finding her opinion utterly insane would justify us in regarding her as an epistemic inferior. However, as Elga points out, if the situation is symmetric, that is if our supposed peer has exactly the same doxastic attitude towards our belief that we have towards hers, EW again requires both of us to suspend judgement.³⁷ This last claim, among other things, is wrong. Before focusing on the errors in Elga's argument, however, let's first see what the correct account looks like.

Although we should not normally use our discussion partner's belief to conclude anything about her peer status, extreme cases such as *THE AWKWARD COLLEAGUE* are an exception because, in such cases, her belief functions as a defeater of whatever evidence we may have regarding her peer status. In *THE AWKWARD COLLEAGUE*, Jane seemed to me a pretty smart person, but whatever reasons I had for this belief are undermined by the fact that she thinks that $2+2=5$. Could one really call someone smart who thinks that $2+2=5$? I do not know; or,

³⁷ Elga, "Reflection," 491.

rather, I think that our very understanding of smartness, or intelligence, or competence, is simply not apt to answer that question. Learning that Jane thinks that $2+2=5$ makes it impossible for me to compare her respective degrees of smartness, intelligence, and competence to mine. In order to do so without just begging the question and assuming that I am right about $2+2$, I would have to set aside beliefs that are so fundamental that there is not enough conceptual knowledge left to say what smartness, intelligence, competence, etc. actually are. If some disagreeing interlocutor's thinking differs too radically from ours, we cannot apply our qualitative concepts to it without thereby taking sides.³⁸

In scenarios such as THE AWKWARD COLLEAGUE, there is therefore no point in comparing my probability of being right with Jane's. Our respective peer statuses are *incommensurable*. Thus, I cannot just regard Jane as an epistemic inferior; I rather cannot say anything about whether she is my superior, inferior, or equal, and that is not because I lack some relevant information. There could be no information that would help in this case. As a consequence, EW cannot be applied in cases like THE AWKWARD COLLEAGUE, because EW tells us only how much weight we should give our peers' opinions, but not how much weight we should give to the opinions of people whose peer statuses are incommensurable to ours. Therefore, nothing forces us to revise our beliefs in extreme cases of disagreement. (Since there is obviously no sharp line that divides cases like THE AWKWARD COLLEAGUE from those in which we are not left stranded in assessing competence, we cannot always tell whether EW should be applied.)

Admittedly, this might appear unsatisfying.³⁹ The reason is that we can conceive of a generalisation of EW—call it EW*—which holds that we should also give equal weight to the opinions of those whose peer statuses are incommensurable to ours. Then my disagreement with Jane would indeed force me to suspend judgement on whether $2+2=4$. In favour of EW*, one could argue that we have no reason to take our actual conception of rationality to be right. The consequence would be relativism about rationality, according to which a judgement *p* has to be understood as elliptical for a judgement of the form *p relative to the conception of rationality R*—in the same way as Einstein's relativistic conception of mass entails that judgements like *x has mass M* have to be interpreted as *x has mass M relative to spatio-temporal framework S*, or moral

³⁸ Analogous points might be made for other list definition factors, e.g. well-informedness. It appears, however, that uncertainty concerning the application of our terms for those factors goes hand in hand with uncertainty concerning the application of evaluative terms for characteristics such as competence or rationality, so it is the latter we should primarily focus on.

³⁹ I am grateful to an anonymous reviewer for pressing me on this point.

relativism entails that judgements like *It is morally wrong to ϕ* have to be interpreted as judgements like *It is morally wrong to ϕ relative to moral framework M* .⁴⁰ Against EW*, one has to adduce that relativism about rationality entails what could be called epistemic antirealism about truth, namely that there is no proposition of which we could know that it is objectively true (and not just true relative to some specific conception of rationality). The reason is that every judgement is contestable by someone whose underlying conception of rationality is sufficiently different, so that, according to EW*, disagreement would force us to suspend judgement on whether it is true. (It is of little help to point out that we quite rarely, if at all, encounter extreme cases of disagreement, and that we are therefore almost never forced to suspend judgement on propositions of whose truth we are highly confident. For one thing, it would be odd if our justification to believe such propositions as $2+2=4$ depended on the contingent fact of whether or not we have met someone who honestly denies their truth; and for another thing, it has been argued that merely possible disagreements are of epistemic significance as well,⁴¹ so that mere contestability might already suffice for suspension of judgement.) Hence, there either are no propositions that are objectively true, or, if there are any, we will never be able to establish their truth. In short, EW* is on the one hand more demanding as EW, since it requires a non-trivial further assumption, and comes on the other hand with considerable theoretical costs regarding the nature and epistemology of true propositions. I therefore ignore EW* in what follows, and continue to discuss EW, which, as we have seen, cannot be applied to cases like THE AWKWARD COLLEAGUE, in which the protagonists' peer statuses are incommensurable.

Regarding Elga's way of dealing with extreme cases of disagreement, this means that he is wrong both in thinking that EW requires us to suspend judgement if the extreme case is symmetric (like THE AWKWARD COLLEAGUE) and in thinking that we should regard our discussion partner as epistemically inferior if she does not find our belief as irrational as we find hers. Since in both symmetric and asymmetric cases we cannot reasonably compare the other one's peer status with ours, EW allows us to disregard his or her opinion in both kinds of cases—though in each case not on the basis of epistemic inferiority.

⁴⁰ See Gilbert Harman, "Moral Relativism," in *Moral Relativism and Moral Objectivity*, Gilbert Harman and Judith Thomson (Oxford: Blackwell, 1996), for a comparison between relativism about mass and relativism about morals (1–3, 13, 18–19, 41) as well as for an outline of how moral relativists can explain actual disagreements (32–44). This outline would by and large carry over to the case of relativism about rationality.

⁴¹ See e.g. Marc Andree Weber, "Armchair Disagreement," *Metaphilosophy* 48 (2017).

Recall from section 3 that Elga gives two reasons for revising EW. One is that we have to take into account the specific circumstances of the disagreement (e.g. hot weather). I argued that this should rather be done by relativising peerhood to time. The other is that we should have a plausible way of dealing with extreme cases of disagreement. We now see, however, that the potential occurrence of such cases does not put us under any pressure to refine EW in the way Elga suggests: since the parties to an extreme disagreement cannot reasonably be regarded as peers, EW, in the sense of definition 2, simply does not apply. Thus Elga's refined version of EW remains unmotivated.

7. Conclusion

The two main insights in this paper are, first, and anticipated in a footnote of Elga's, that a likelihood definition—preferably my definition 1—captures the essence of epistemic peerhood better than any kind of list definition and should therefore be adopted; and, second, that this adoption necessitates a shift of focus in the debate on peer disagreement: the alleged theoretical alternatives to EW, which are in fact compatible with it, should not be taken to entail that we should give our peers' beliefs less than equal weight, but only that even very competent and well-informed people may easily fail to be our peers. Revised accordingly, however, these alternatives might appear harder to believe.

Further insights in this paper include criticisms of several of Elga's views that are based on or related to his understanding of epistemic peerhood (the bootstrapping argument; his version of EW; the thesis that EW does not hold in clusters of controversy; his way of dealing with extreme cases of disagreement). They also include the idea that in extreme cases of disagreement—but only in those—we lack the common ground needed to compare people's likelihoods of being right. This incommensurability, however, sets only somewhat inextensive limits to our ability to ascribe peerhood and apply EW. As you will never in your life encounter a person like my colleague Jane, you should always give your peers' opinions equal weight.

DISCUSSION NOTES/ DEBATE

GETTING GETTIER RIGHT: REPLY TO MIZRAHI

Philip ATKINS

ABSTRACT: Moti Mizrahi has argued that Gettier cases are misleading, since they involve a certain kind of semantic failure. In a recent paper, I criticized Mizrahi's argument. Mizrahi has since responded. This is a response to his response.

KEYWORDS: analysis of knowledge, Gettier cases, semantic reference, speaker's reference, ambiguous designators, intuition mongering

Moti Mizrahi has argued that Gettier cases are misleading, since they involve a certain kind of semantic failure. In a recent paper, I criticized Mizrahi's argument. Mizrahi has since responded. This is a response to his response.¹

Mizrahi begins his response to me by mentioning some things that he finds peculiar about my critique. I said that Gettier's original two cases are genuine counterexamples to the Justified True Belief analysis of knowledge (henceforth, the JTB analysis). And yet, in replying to Mizrahi's original paper, I revised Gettier's first case. Mizrahi asserts that if I needed to revise this case, then it is not a genuine counterexample. Otherwise why would I need to revise it?

Let me explain how I understand this matter. Gettier cases standardly elicit the intuition that the relevant agent lacks knowledge even though the agent has a justified true belief. If this intuition is accurate, then Gettier cases are genuine counterexamples to the JTB analysis. And when I say that Gettier cases are genuine counterexamples to the JTB analysis, I mean only that they are cases in which the relevant agent has a justified true belief and yet lacks knowledge. Now, Mizrahi tries to call into question the accuracy of the Gettier intuition by arguing that it may result, not from any epistemic failure, but rather from a certain kind of semantic failure. As Mizrahi put it in his original paper, we who have the intuition "may simply be mistaking semantic facts for epistemic facts when we

¹ For the original Gettier cases, see Edmund Gettier, "Is Justified True Belief Knowledge?" *Analysis* 23 (1963): 121-123. For Mizrahi's first paper, see "Why Gettier Cases Are Misleading," *Logos & Episteme* 7 (2016): 31-44. For my response, see "Are Gettier Cases Misleading?" *Logos & Episteme* 7 (2016): 379-384. For Mizrahi's response to my response, see "Why Gettier Cases Are Still Misleading," *Logos & Episteme* 8 (2017): 129-139.

consider Gettier cases.”² If this alternative explanation of the Gettier intuition is correct, then the intuition should be absent when considering Gettier cases where it is clear that there is no such semantic failure. In my response to Mizrahi, I tried to provide such a case by simply tweaking one of Gettier’s original examples, and I then observed that, when considering this case, there remains the intuition that the relevant agent lacks knowledge. This means that Mizrahi’s alternative explanation of the Gettier intuition is incorrect. So, even though I revised Gettier’s first case in replying to Mizrahi, I continue to hold that it is a genuine counterexample to the JTB analysis. I find nothing peculiar about this position.³

In his preliminary remarks, Mizrahi also says that I presented “a somewhat inaccurate picture of the state of the debate over the status of Gettier cases as a ‘refutation’ of the JTB analysis of knowledge.”⁴ This is because, according to Mizrahi, “epistemologists have long recognized that Gettier’s original cases are problematic,”⁵ since these cases involve inferences from false premises, or “false lemmas.” I deny the charges. It is true that epistemologists have long recognized that Gettier’s original cases involve inferences from false lemmas. But it is misleading to say that these epistemologists took Gettier’s original cases to be “problematic,” for these epistemologists generally took Gettier to have been successful in refuting the JTB analysis. In light of Gettier’s counterexamples, some epistemologists were moved to argue that knowledge is actually justified true belief *without false lemmas*,⁶ while others claimed that even this theory succumbs

² “Why Gettier Cases Are Misleading,” 33.

³ Consider an analogy. Suppose that someone tries to explain away my intuition by pointing to the fact that the relevant agent in Gettier’s case is named ‘Smith.’ Since I was recently dumped by someone named ‘Smith,’ I may be holding a grudge and for this reason I may be unwilling to ascribe knowledge to anyone named ‘Smith.’ To rebut this alternative explanation of my intuition, I need only revise the case so that the relevant agent is named something else, and then observe that the agent still seems to lack knowledge. And even though I have revised Gettier’s original case to rebut the alternative explanation, I can coherently hold that Gettier’s original case is itself a counterexample to the JTB analysis.

⁴ “Why Gettier Cases Are Still Misleading,” 130.

⁵ “Why Gettier Cases Are Still Misleading,” 129.

⁶ Michael Clark “Knowledge and Grounds,” *Analysis* 24 (1963): 46-48. David Armstrong went further and suggested that one’s lemmas must not only be true, but known to be true. See his *Belief, Truth and Knowledge* (Cambridge University Press, 1974), 152-154. Both agree that Gettier’s original cases are counterexamples to the JTB analysis. I am aware of very few philosophers who have questioned whether Gettier’s original cases are genuine counterexamples. There is, for example, Joseph Margolis, “The Problem of Justified Belief,” *Philosophical Studies* 23 (1972): 405-409; and Meyers and Stern, “Knowledge Without Paradox,” *Journal of Philosophy* 52 (1973): 147-160. But, unlike Mizrahi, these philosophers readily accept

to Gettier-style counterexamples.⁷ Mizrahi questions why I would neglect to mention these things in my reply to him. The answer is that I was defending the claim that Gettier's original cases refute the JTB analysis, not the claim that they refute the stronger theory that knowledge is justified true belief without false lemmas. To my knowledge, nobody has ever held that Gettier's original cases refute this stronger theory and I am not sure why Mizrahi thinks that I should have discussed it.

A note on terminology. Mizrahi thinks that "refutations" of theories are "conclusive proofs" that those theories are false. When I say that Gettier refuted the JTB analysis, I mean only that he provided counterexamples to the JTB analysis (that is, cases in which the relevant agent has a justified true belief and yet lacks knowledge). This seems somewhat different from saying that Gettier provided a "conclusive proof," since it is always possible that someone will try to explain away our intuitions, which is exactly what Mizrahi tries to do to the Gettier intuition. I will gladly allow that this point is merely terminological. Whenever I say that Gettier refuted the JTB analysis, I can be understood to mean that Gettier's original cases are counterexamples to the JTB analysis. However, there are more substantive differences between Mizrahi and me. He is generally skeptical of appeals to intuitions. In fact, he accuses me of "mere intuition mongering," though it is not clear to me exactly what he means.⁸ My own modest view is that one's intuitions count as good evidence for or against philosophical theories, but they are also defeasible, as all forms of evidence are defeasible. Various things might undermine the evidential weight of an intuition, such as the fact that many others lack the intuition, or a plausible alternative account of why one has the intuition. As I have made clear, I think that Mizrahi's alternative account of the Gettier intuition, whereby it results from some mistake that we are making regarding epistemic facts and semantics facts, does not successfully explain the Gettier intuition. And, though I am sure that there are those who do not have

that the agents in Gettier's cases lack knowledge. Instead, they deny that the agents have formed justified beliefs. Suffice it to say, this position has never been popular.

⁷ See, for example, Ernest Sosa, "The Analysis of 'Knowledge That P,'" *Analysis* 25 (1964): 1-8; John Turk Saunders and Narayan Champawat, "Mr. Clark's Definition of 'Knowledge,'" *Analysis* 25 (1964), 8-9; Richard Feldman *Epistemology* (Prentice Hall, 2003), 31-33. Of course, the debate continues. For a relatively recent discussion, see Michael Levin "Gettier Cases Without False Lemmas?" *Erkenntnis* 64 (2006): 381-391.

⁸ Mizrahi discusses the topic in "Intuition Mongering," *The Reasoner* 6 (2012): 169-170; and also "More Intuition Mongering," *The Reasoner* 7 (2013): 5-6. He does not provide a definition of 'intuition mongering,' so I am left to conclude that he means nothing more than 'relying on our intuitions,' or alternatively 'relying on how things seem to us.'

the Gettier intuition, recent empirical research indicates that it is widely shared across cultures.⁹

Speaking of empirical research, though Mizrahi does not emphasize experimental philosophy in his original paper, he mentions it in his reply to me as evidence that appeals to intuition are “rather controversial.”¹⁰ This despite the fact that recent research indicates that the Gettier intuition is remarkably pervasive, as I mentioned above. In any event, Mizrahi is free to argue against Gettier on the general grounds that philosophical intuitions are unreliable, as suggested, perhaps, by empirical research. But this is not the main argument that he made in his original paper, which was much narrower, and which was the argument that I wanted to rebut. As Mizrahi himself wrote, the main argument of that paper did “not depend on experimental results concerning Gettier intuitions.”¹¹

Before discussing the details of Gettier’s original cases, there is one more preliminary point worth discussing. Mizrahi claims that I denied that Gettier’s original cases involve ambiguous designators. Though I do not believe that ambiguous designators are responsible for the Gettier intuition, I never denied that Gettier’s original cases involve ambiguous designators. Indeed, I purposefully avoided using the expression ‘ambiguous designator’ in my initial reply to Mizrahi. This is because I felt that Mizrahi was using the term in an idiosyncratic way. In fact, though Mizrahi relies on Saul Kripke’s distinction between semantic reference and speaker’s reference,¹² Mizrahi uses ‘ambiguous designator’ in a way that Kripke would not. Remember that Kripke was concerned with the claim that definite descriptions, such as ‘the man drinking champagne,’ are ambiguous, since they can be used attributively or referentially.¹³ Kripke denies that ‘the man drinking champagne’ is ambiguous, at least in the way that ‘bank’ is ambiguous.¹⁴ He argues that the distinction between attributive use and referential use is an instance of a more general distinction, between semantic reference and speaker’s

⁹ Edouard Machery, Stephen Stich, David Rose, Amita Chatterjee, Kaori Karasawa, Noel Struchiner, Smita Sirker, Naoki Usui, and Takaaki Hashimoto, “Gettier Across Cultures” *Noûs* (forthcoming).

¹⁰ “Why Gettier Cases Are Still Misleading,” 131.

¹¹ “Why Gettier Cases Are Misleading,” 32, fn. 4.

¹² Saul Kripke, “Speaker’s Reference and Semantic Reference,” *Midwest Studies in Philosophy* 2 (1977): 255-276.

¹³ Keith Donnellan, “Reference and Definite Descriptions,” *The Philosophical Review* 75 (1966): 281-304.

¹⁴ In general, Kripke is wary of invoking ambiguities in philosophical debates. He writes, “It is very much the lazy man’s approach in philosophy to posit ambiguities when in trouble,” 268. Thanks to Matt Griffin for discussing this point with me.

reference. This distinction applies even to proper names, such as ‘Barack Obama.’ The semantic referent of the name is Barack Obama, the former president, but the speaker’s referent, on a certain occasion of use, might be someone else entirely. And yet, even if there is a divergence between semantic reference and speaker’s reference, it would be inappropriate to conclude that ‘Barack Obama’ is ambiguous. This, anyway, is Kripke’s contention. Mizrahi would say, apparently, that the proper name is ambiguous, since he thinks that divergences between semantic reference and speaker’s reference are indicative of ambiguous designators.

I side with Kripke in thinking that ‘Barack Obama’ is unambiguous. After all, if we say that ‘Barack Obama’ is ambiguous, then we will have to say that *every* singular expression is ambiguous, at least potentially, since every singular expression admits of divergences between speaker’s reference and semantic reference. I am not willing to countenance such an explosion of ambiguities.¹⁵ Ultimately, however, I am inclined to think that the matter is not of central importance. Mizrahi’s main argument can be stated without explicitly mentioning ambiguous designators. That argument is that Gettier cases involve a certain kind of semantic failure, where the semantic referent of a term is different from the speaker’s referent, and that this semantic failure is plausibly responsible for our intuitions regarding Gettier cases. So, according to Mizrahi, those cases are misleading.

Now, let us consider more closely Gettier’s first case. Smith comes to have strong evidence for believing that Jones is the man who will get the job and that Jones has ten coins in his pocket. Smith then infers

(I) The man who will get the job has ten coins in his pocket.

This is true, not because Jones has ten coins in his pocket, but because Smith has ten coins in his pocket and Smith happens to be the man who will get the job. Mizrahi argued in his original paper that ‘coins’ is an ambiguous designator, since the speaker’s referent is the set of coins in Jones’s pocket, which is not the semantic referent of ‘coins.’ I argued that, even if there is a difference here between the speaker’s referent and the semantic referent, it does not plausibly account for the intuition that Smith lacks knowledge. For the case can be easily revised so that Smith instead comes to infer

(I*) The man who will get the job is handsome.

¹⁵ I imagine that Mizrahi would be fine with this explosion of ambiguities. In fact, as we will see, Mizrahi is willing to say that even complex quantificational expressions, such as ‘there is someone who,’ are ambiguous.

It seems that Smith fails to know (I*), and yet it is unreasonable to insist that there is a divergence between the speaker's referent of 'handsome' and the semantic referent of 'handsome.' I considered the possibility that 'the man who will get the job' is such that there is a divergence between the speaker's referent and the semantic referent. Even though Mizrahi did not focus on the definite description, and would not be the first to argue that there is a divergence there between speaker's reference and semantic reference,¹⁶ I noted that the case can again be revised so that Smith infers

(I**) There is someone who is getting a job and handsome.

Intuitively, Smith fails to know (I**), and yet there is no definite description whatsoever.¹⁷ So, even if Gettier's first case involves an ambiguous designator, in Mizrahi's special sense, this fact would not successfully explain away our intuition about the case.

In Mizrahi's response to me, he does not share his intuitive judgment of these cases. He *does* suggest that the intuition that Smith lacks knowledge results from my having been taught that this intuition is the "right" one. It is not clear to me how to respond to such speculation, except to say that, upon serious reflection, I still judge that Smith lacks knowledge in these cases (I am quite capable of thinking for myself). Mizrahi also suggests that the intuition results from the fact that these cases involve false lemmas. Indeed, this is a plausible account of the intuition, insofar as it is a plausible account of why Smith lacks knowledge in these cases. In other words, it is plausible that knowledge requires more than justified true belief. Perhaps it requires the absence of false lemmas.¹⁸ But Gettier never argued otherwise. To repeat what I said above, I am not aware of anyone

¹⁶ Mizrahi himself cited Adrian Heathcote "Truthmaking and the Gettier Problem," in *Aspects of Knowing: Epistemological Essays*, ed. Stephen Hetherington (Oxford: Elsevier, 2006), 151-168; and also Christoph Schmidt-Petri "Is Gettier's First Examples Flawed?" in *Knowledge and Belief*, eds. Winfried Löffler and Paul Weingartner (Kirchberg am Wechsel: ALWS, 2003), 317-319.

¹⁷ I should have mentioned in my earlier paper that this version of Gettier's case is basically the same as Keith Lehrer's Nogot/Havit case. In that case, I have strong evidence for believing that Mr. Nogot, who is in my office, owns a Ford. I then infer that someone in my office owns a Ford. It turns out that someone in my office *does* own a Ford, but, unbeknownst to me, it is Mr. Havit rather than Mr. Nogot. Here I have a justified true belief that someone in my office owns a Ford, but this belief does not count as knowledge. See "Knowledge, Truth and Evidence," *Analysis* 25 (1965): 168-175. Lehrer's case is probably cleaner than mine, since mine requires some minor qualifications regarding the claim that Smith fails to know (I**). See fn. 9 of "Are Gettier Cases Misleading?"

¹⁸ See fn. 6 above.

who has thought that Gettier's original cases refute the theory that knowledge is justified true belief *without false lemmas*. What epistemologists have standardly held is that Gettier's original cases refute the theory that knowledge is justified true belief. This is the position I want to defend against Mizrahi. The stronger theory of knowledge, whereby knowledge requires the absence of false lemmas, is irrelevant here.

Mizrahi proceeds to say that 'someone,' or more fully 'there is someone who,' is ambiguous in (I**). He thinks that the semantic referent is Smith, while the speaker's referent is Jones. It is odd to suggest that the semantic referent of the complex quantificational expression 'there is someone who' is one particular man, but set that issue aside. Mizrahi's main point is that the content of Smith's belief can be interpreted in two ways.

1. There is someone (=Smith) who is getting a job and handsome.
2. There is someone (=Jones) who is getting a job and handsome.

Mizrahi notes that (2) is false and (1) is not believed by Smith. But neither (1) nor (2) is the relevant interpretation of what Smith believes after inferring (I**). The content of Smith's belief is, quite simply, the general proposition *that there is someone who is getting a job and handsome*, not the singular proposition *that there is someone, specifically Smith, who is getting a job and handsome*, nor the singular proposition *that there is someone, specifically Jones, who is getting a job and handsome*.¹⁹ The question is whether *that belief*, whose content includes neither Smith nor Jones, counts as knowledge. Intuitively, the answer is that it does not. Mizrahi might insist that Smith does not actually believe the fully general proposition *that there is someone who is getting a job and handsome*, but this maneuver would be desperate. Why should Smith be cognitively unable to form such a rudimentary belief?

Mizrahi makes one other curious claim about my revision of Gettier's first case. He thinks that Smith must be reasoning as follows:

- a. Jones is getting the job.
- b. Therefore, there is someone who is getting the job.
- c. Jones is handsome.
- d. Therefore, there is someone who is handsome.

¹⁹ Similarly, in the Nogot/Havit case, the content of my belief is understood to be the general proposition *that someone in my office owns a Ford*, not the proposition *that someone in my office, specifically Nogot, owns a Ford*. See fn. 17 above.

Philip Atkins

e. Therefore, there is someone who is getting the job *and* there is someone who is handsome.

Mizrahi asserts that “Smith’s evidence supports (e), not the belief that the one who will get the job and the one who is handsome *are one and the same person*.”²⁰ But these remarks are incorrect. To be as explicit as possible, Smith is actually reasoning as follows:

a*. Jones is getting the job.

b*. Jones is handsome.

c*. Therefore, Jones is getting the job and handsome.

d.* Therefore, there is someone who is getting the job and handsome.

Formally, Smith’s line of reasoning could be spelled out as follows:

Gj

Hj

$\therefore Gj \bullet Hj$

$\therefore \exists x[Gx \bullet Hx]$.

Now, it is stipulated that Jones has strong evidence for believing both (a*) and (b*). From (a*) and (b*) he infers (c*), and from (c*) he infers (d*). Each inference is valid and (d*) clearly entails that the one who is getting the job and the one who is handsome are identical.

As for Gettier’s second case, I observed in my initial reply that Mizrahi misinterpreted it, which undermined his case for semantic failure. In that case, Smith has strong evidence for believing that Jones owns a Ford. His evidence is that “Jones has at all time in the past owned a car, and always a Ford, and that Jones has just offered Smith a ride while driving a Ford.”²¹ Smith makes a rudimentary logical inference and says the following:

(h) Either Jones owns a Ford or Brown is in Barcelona.

It turns out that (h) is true, not because Jones owns a Ford, but because Brown is in Barcelona. Mizrahi presented the case so that Smith infers (h) from

(g) Either Jones owns a Ford or Brown is in Boston,

which was inferred from

(f) Jones owns a Ford.

²⁰ “Why Gettier Cases Are Still Misleading,” 135.

²¹ “Is Justified True Belief Knowledge?” 122.

In fact, Smith infers both (g) and (h) directly from (f). Moreover, Mizrahi presented the case, incorrectly, as involving two separate men, one of whom is the speaker's referent of 'Jones' and one of whom is the semantic referent of 'Jones.' As far as I can tell, Mizrahi does not address these mistakes in his reply to me. He does, however, accuse *me* of failing to get Gettier right. In his reply, he presents Smith's reasoning as follows:

- i. Smith has at all times in the past within Smith's memory owned a car.
- ii. Smith has at all times in the past within Smith's memory owned a Ford.
- iii. Jones has just offered Smith a ride while driving a Ford.
- iv. Therefore, Jones owns a Ford.
- v. Therefore, either Jones owns a Ford or Brown is in Barcelona.

Mizrahi then writes: "Contrary to what Atkins suggests, Smith cannot simply make 'a rudimentary logical inference' from (i)-(iii) to (v), since (v) does not follow from (i)-(iii). Rather, (i)-(iii) are evidence for (iv), and then Smith infers (v) from (iv) by a 'rudimentary logical inference,' namely, addition."²² Well, I agree that (i)-(iii) are Smith's evidence for (iv) and that Smith, after concluding that (iv) is true, infers (v) from (iv). However, I did not say otherwise. As for whether I "suggested" otherwise, I will have to let the reader decide. My presentation of Gettier's case was essentially the same as the presentation that I give above.

I have little space to discuss the other things that Mizrahi says about Gettier's second case, so I will only make a few brief points. While reflecting on the nature of time, Mizrahi observes that a child could not correctly reason as follows:

Barack Obama has at all times in the past within my memory been the US president.

Therefore, Barack Obama is the US president at present (where the present time is January 21, 2017).

"The problem," according to Mizrahi, "is that 'Barack Obama' is referentially ambiguous in this context."²³ Philosophers of language would tend to disagree. The problem is not the proper name, but rather the definite description 'the US president,' which designates different individuals with respect to different times. Regardless, Mizrahi proceeds to argue that 'Jones' is also referentially ambiguous:

²² "Why Gettier Cases Are Still Misleading," 136.

²³ "Why Gettier Cases Are Still Misleading," 137.

“the reference of ‘Jones’ in (i)-(iii) was fixed at some particular time in the *past*, since (i)-(iii) are based on what Smith *remembers* about Jones, whereas ‘Jones’ in (iv) is supposed to pick out the *present* Ford owner. This switch in reference... makes Gettier’s Case II appear like a genuine counterexample to JTB, even though it is not.”²⁴

So, ‘Jones’ supposedly undergoes a shift in reference, designating one individual and then later designating a different individual. Who specifically are these individuals? The initial referent must be Jones, the man whom Smith remembers as having owned a Ford. But who is the other individual? Mizrahi suggests that it is “the *present* Ford owner.” But if ‘Jones’ designates a Ford owner in (iv), then (iv) would be true, whereas in fact (iv) is stipulated to be false. Perhaps it is the case that ‘Jones’ initially designates *past* Jones and then comes to designate *present* Jones. This is perhaps why Mizrahi writes in his conclusion that “Smith has *past* Jones in mind, for Smith’s evidence is about *past* Jones, not about *present* Jones.”²⁵ At the risk of wading into murky metaphysical waters, I submit that Smith has both past Jones and present Jones in mind, and that his evidence is about both past Jones and present Jones, seeing as how past Jones and present Jones *are the same person*. And if we agree that there is only one individual here, not two separate individuals, then we cannot coherently say that there is a switch in reference.

In any event, Mizrahi’s alternative explanation of the Gettier intuition falls apart if we simply imagine that Jones is currently in the company of Smith. Instead of thinking (iv), based on memory, he comes to the following conclusion, based on direct perception:

(iv*) Jones is standing in front of me.

And then, from (iv*), Smith validly infers

(v*) Either Jones is standing in front of me or Brown is in Barcelona.

Of course, (iv*) is false. Jones is standing behind Smith, but appears to be standing in front of Smith due to some cleverly placed mirrors. Nonetheless, by sheer coincidence, Brown is in Barcelona and so (v*) is true. Surely anyone who agreed with Gettier’s conclusion regarding (v) will come to the same conclusion regarding (v*). Specifically, they will come to the conclusion that Smith fails to know (v*), even though Smith is justified in believing (v*) and it turns out that (v*) is true. And yet here there can be no quibbles about past Jones and present Jones.

²⁴ “Why Gettier Cases Are Misleading,” 138.

²⁵ *Ibid.*

There can be no talk of 'Jones' somehow changing its reference over time. It stands to reason that the considerations adduced by Mizrahi are not responsible for our intuitions about Gettier's second case. A much better account of our intuitions would be that knowledge is not justified true belief.

AGAINST SENSE-DATA AS STRUCTURED UNIVERSALS

Landon D. C. ELKIND

ABSTRACT: I critically discuss a new proposal for a metaphysics of sense-data. This proposal is due to Peter Forrest. Forrest argues that, if we accept Platonism about universals, sense-data are best understood as structured universals—in particular, as structured universals with temporal and spatial properties as components. Against this proposal, I argue sense-data as structured universals are not universals at all.

KEYWORDS: sense-data, perception, structured universals, universals, metaphysics

1. A New Metaphysics of Sense-Data

I criticize a revisionist conception of sense-data proposed by Peter Forrest in 2005. For Forrest, the object of perception (the sense-datum) is one or more structured universals.¹ I explain the ‘structured’ qualifier shortly. But sense-data theories are wildly unpopular. In a 2009 survey just 3.1% of philosophers reported accepting the sense-data theory.² So why should you care about an in-house debate among sense-data theorists?

Here are two reasons. First, whether you find a sense-data theory remotely plausible depends on what a sense-data theory is. Yet Forrest's sense-data theory is in important ways a radical departure from the standard conception of a sense-data theory. If his proposal is viable, then our definition of the sense-data theory requires correction.

Second, while sense-data have sunk into ill-repute, Platonism remains the most popular metaphysics of abstract objects among respondents to that same 2009 survey.³ In this context, by ‘Platonism’ I mean “belief in possibly uninstantiated universals.”⁴ Now Forrest offers to cleanly solve the problem of perception if we

¹ Peter Forrest, “Universals as Sense-Data,” *Philosophy and Phenomenological Research* LXXI, 3 (2005): 622.

² David Bourget and David J. Chalmers, “What Do Philosophers Believe?” *Philosophical Studies* 170 (2014): 476.

³ Bourget and Chalmers, “What Do Philosophers,” 475.

⁴ Forrest, “Universals as Sense-Data,” 622.

just grant him Platonism.⁵ So if we are among the 39.3% that accept Platonism, or among the 3.1% that sorely need a more attractive version of the sense-data theory, we should seriously consider Forrest's proposal.

Let us return to the issue of the definition of the sense-data theory. Howard Robinson defines the core of the sense-data theory with two claims:⁶

1. Perceptual experience is *relational*: it is analyzable into an act and an object.
2. The object of perception (the sense-datum), is
 - a. an object of perceptual awareness,
 - b. non-physical,
 - c. private to some subject,
 - d. what possesses sensible qualities, and
 - e. not (merely or wholly) intentional or representation.

Let us see how Forrest's new metaphysics of sense-data radically departs from this conception of sense-data in (2). On his view the sense-datum is a "complex relation;" it is complex because it involves both a quality and conditions of perception. Perception of a tree, say, is a perception of a structured universal *being-tree-like-and-in-front-of-the-subject*.⁷

So far time has not surfaced, but Forrest brings the view to completion by including temporal relations as well as spatial ones.⁸ Thus Forrest holds sense-data are structured universals, namely, spatiotemporally-structured universals like *being-red-and-in-front-of-the-subject-and-present-to-the-subject*: a thus-and-there-and-then.⁹

These additional conditions make the object of perception a structured universal rather than a simple universal like *redness*. Call such sense-data (that is,

⁵ "If Platonism is otherwise acceptable we'd be crazy not to identify sense-data with universals." Forrest, "Universals as Sense-Data," 631.

⁶ Howard Robinson, *Perception* (New York: Routledge, 1994), 1-2.

⁷ Forrest, "Universals as Sense-Data," 622-623.

⁸ "Where in this paper I have used purely spatial relation in the analysis of perception we should replace them by spatio-temporal ones." Forrest, "Universals as Sense-Data," 629.

⁹ Forrest, "Universals as Sense-Data," 629. For Forrest, not *all* modes of sensory awareness include a temporal component: imagination involves a structured universal with a spatial location but without a temporal location; in contrast, memory involves a structured universal with a spatial location and an temporal location "earlier than" the time of the imagining. The temporal component in a sense datum, the *then* component, can be leveraged to distinguish memory, imagination, and other forms of sensory awareness; further, the temporal component "explains the vividness of perception as due to its present tense character." Forrest, "Universals as Sense-Data," 630.

qualities structured by spatiotemporal conditions of perception) “atomic,” as Forrest holds such sense-data admit of no further analysis.¹⁰

Note how sense-data, so understood, differ from traditional sense-data. Sense-data as structured universals are neither mental (as on most theories) nor physical, but abstract entities. And they are not obviously private: just as we can both be aware of the same universal *redness*, so too we can both be aware of the same universal *being-red-and-in-front-of-the-subject-and-present-to-the-subject*. Sense-data as structured universals also do not possess sensible qualities but are partially constituted by them: sensible qualities are constituents of a (structured universal) sense-datum instead of, as on the traditional version, inhering in (particular) sense-data. So (2)(b)-(d) above might need to be revised.

This new metaphysics of sense-data comes with many benefits. It offers a straightforward analysis of hallucination.¹¹ It solves two long-standing problems with sense-data views.¹² It solves the problem of indeterminacy in objects of perception (witness the well-worn speckled hen).¹³ And it solves the problem posed by the apparent seamless transition between hallucinatory perception and veridical perception (aptly described by Mark Johnston).¹⁴ This new sense-data theory looks much more attractive (to Platonists) than the traditional theory, or at least the new theory avoids the usual objections.

2. Sense-Data as Structured ‘Universals’

Problematically, this new metaphysics of sense-data is internally incoherent. For sense-data as structured universals are not repeatable, and so are not really universals. Note that being repeatable is a condition thought by many to be necessary for being a universal. Borrowing our wording from Lord Russell, a universal is distinguished by its potentially being multiply located; particulars, in contrast, cannot be multiply located.¹⁵

Now sense-data as structured universals, namely, as spatially and temporally located qualities, lack this distinguishing mark. For a sense datum being partly

¹⁰ “Atomic sensations have structure; they consist of being appeared to thus-and-here...[which] is not the same as being appeared to thus and being appeared to here.” Forrest, “Universals as Sense-Data,” 622.

¹¹ Forrest, “Universals as Sense-Data,” 623.

¹² Forrest, “Universals as Sense-Data,” 623-624.

¹³ Roderick Chisholm, “The Problem of the Speckled Hen,” *Mind* 51, 204 (1942): 368-373, 368.

¹⁴ Mark Johnston, “The Obscure Object of Hallucination,” *Philosophical Studies* 120, 1/3 (2004): 113-183, 122-123.

¹⁵ Bertrand Russell, *The Problems of Philosophy* (New York and London: Home University Library, 1912), 145.

constituted by its a spatial and temporal location implies the impossibility of its being multiply located. A quality q located at l at time t (write ' (q, l, t) ') is not identical with (q, l', t') if the locations and times differ; that is, if $l \neq l'$ and $t \neq t'$, then $(q, l, t) \neq (q, l', t')$. So sense-data as structured universals, to be properly called universals, must be potentially multiply located unless some non-standard distinguishing characteristic of universals is on offer. As sense-data as structured universals cannot be multiply located, they are universals in name only.

There are at least three repairs available. (1) One could claim that structured universals may be multiply located because symbols for space and time locations are *indexical*, like 'here' and 'now.' Then the current object of sensory awareness is a quality q in a here-location h at a now-time n where the location h and time n vary in referent with occasions of use. So even if one sees a quality q at distinct locations h and h' and distinct times n and n' , we might still have $(q, h, n) = (q, h', n')$ because the quality q is 'here-and-now' in both cases. Yet this does not solve the problem. For 'now' means something like 'at the speaker's present time,' and similarly for 'here.' And though one can see a tree in many heres-and-nows, the meaning of 'here-and-now' differs on each occasion of the subject's perceiving. It is only by equivocating over the meaning of 'here-and-now' that one could claim that $(q, h, n) = (q, h', n')$. So we still lack a universal: in each instance of perception, the meaning of here-and-now remains a particular space-time.

(2) One can claim that all sense-data as structured universals have determinable locations or (inclusive) times, not determinate ones. So no sense-data would occupy a determinate spatial and temporal location but would be repeatable within a determinable range. But sense-data in at least some cases have a determinate location, not an indeterminate one. The sense-datum of my pen, say, is in a determinate location for a determinate (stretch of) time. So this proposal sorely needs as a supplement some account of when one's sense-datum has a determinate location and when it does not. I doubt a principled account is available here.

(3) One might revise the view so that sense-data have only a spatial component and no temporal component.¹⁶ One could then claim that the temporal component is a feature of one's perceptual awareness rather than a feature of the sense-datum. That is, a perceiving subject is aware at time t of a spatially located

¹⁶ I do not find it plausible to say that a sense-datum has a temporal component but not a spatial one. First, it is not clear how to relate distinct sense-data, like two visual patches, that one perceives at the same time without appealing to space. Second, being aware of *redness-at-t* is inadequate; one would need to be aware of *an-instance-of-red-at-t*. And this is just a red particular sense-datum described in different words.

sensible quality, that is, a thus-and-there written ' (q, t) '. Then sense-data as structured universals can be multiply instantiated as they are objects of temporally-distinct acts of perceptual awareness.¹⁷

I find three difficulties with this repair. First, it admits objects with a spatial location but no temporal location. This contradicts the current orthodoxy in physics, on which space and time are intimately bound into space-time.¹⁸ Secondly, it introduces an ontological asymmetry between space and time. Platonic universals are constituted outside space and time; more colorfully, they exist "nowhere and nowhen."¹⁹ (In contrast, Aristotelian universals are constituted in space and time.) Sense-data as structured universals are constituted inside space and outside time; being spatially-located qualities without temporal location, they exist 'somewhere and nowhen.' The metaphysical possibility of such objects requires some (independent) justification and explanation. I see no hope for this project.

Note also that the repair sacrifices some of the benefits of the new metaphysics of sense-data. In particular, we lose Forrest's account of memory. So I recommend that we reject the theory of sense-data as structured universals unless someone finds it sufficiently worth saving to offer an account of universals with only spatial components.

In short, repair (3) fixes the incoherence of Forrest's view. Sense-data as structured universals are now genuine universals. But the fix is unprincipled and introduces severe difficulties. It is rather like a premise offered solely because it saves the conclusion. I say we are better off rejecting Forrest's new and revisionary metaphysics of sense-data.²⁰

¹⁷ Forrest, "Universals as Sense-Data," 628-629.

¹⁸ Michael Huemer, *Skepticism and the Veil of Perception* (Lanham: Rowman & Littlefield, 2001), 154. One might reply that sense-data occupy phenomenal space and time, not external world space-time. If defensible, this reply resolves the first objection. It does not answer the remaining two objections.

¹⁹ Russell, *Problems*, 153.

²⁰ I thank my colleagues at the University of Iowa for their feedback on this paper. I am especially grateful to Ali Hasan and to Greg Stoutenburg for the clarification of my thoughts and the abundant pleasure resulting from conversations with them.

THE EPISTEMIC CONSEQUENCES OF FORCED CHOICE

Mark SCHROEDER

ABSTRACT: In “Stakes, Withholding, and Pragmatic Encroachment on Knowledge,” I used a variety of cases, including cases of forced choice, to illustrate my explanation of how and why some pragmatic factors, but not others, can affect whether an agent knows. In his recent contribution, Andy Mueller argues that cases of forced choice actually pose a dilemma for my account. In this paper I reply.

KEYWORDS: knowledge, pragmatic encroachment, Andy Mueller, Pascal’s wager

A number of authors have converged on the consensus that practical factors can affect whether a subject knows – an idea that has come to be known as ‘pragmatic encroachment’ on knowledge. Most theorists who accept this conclusion offer evidence *that* it is true, but few have explored *why* it is true, and most discussions glance quickly over the idea that ‘stakes’ matter without being careful about exactly which kinds of practical factors can matter, and *how* they can matter. But it is important to understand how it could be true that practical factors can make a difference for knowledge. We know that Pascal’s wager – even if it is a good one – does not make it easier to know that God exists. And cases like Pascal’s have made it seem obvious to many philosophers over a long period of time that practical considerations are not the right kind of thing to be relevant in epistemology. So defenders of pragmatic encroachment must explain how practical factors *could* matter, and they must do so in a way that clarifies whether we can accept pragmatic encroachment without being led to the conclusion that Pascal’s wager makes it easier to know that God exists – or similarly, for other pragmatic arguments for belief.

This is why I have been interested, in my work, in explaining how practical considerations do matter for knowledge. A careful and proper way of distinguishing between purely ‘epistemic’ and non-epistemic reasons, I have argued,¹ can rule out Pascalian considerations but still leave room for some practical considerations to count as properly epistemic reasons against belief. I

¹ Mark Schroeder, “What Makes Reasons Sufficient?” *American Philosophical Quarterly* 52, 2 (2015): 159-170.

have argued² that there are, in fact, properly epistemic reasons against belief that are not evidence against the content of that belief, and that among those are considerations whose import is transparently practical – considerations deriving from the consequences of relying on a false belief. I have presented a model of the dynamics of epistemic reasons³ in order to show that the assumption that there really are pragmatic reasons against belief that take this form yields plausible predictions in a wide range of cases, and I have shown⁴ how this model extends to plausible predictions about pragmatic encroachment on knowledge, as well as on rational belief. And in joint work with Jake Ross⁵ I have been developing an account of the nature of outright belief that I am endeavoring to show in work in progress⁶ combines with our best available understanding of the general distinction between right and wrong kinds of reasons to explain *why* costs of relying on a false belief – but not Pascalian considerations in general – are indeed right-kind reasons against belief. Together, I take this body of work to constitute an illuminating and general picture of the *why* and *how* of pragmatic encroachment, and of the positive contributions in all of my work, this set of ideas is the one I am most confident of actually being *true*.

In “Pragmatic or Pascalian Encroachment? A Problem for Schroeder’s Explanation of Pragmatic Encroachment,” Andy Mueller poses an objection to this set of ideas, drawn from one of my own examples, originally described in “Stakes, Withholding, and Pragmatic Encroachment on Knowledge.” In the example, a subject faces a forced choice – she must choose now which of two banks to head towards to try to deposit her check, she knows that only one of them is open now, and she has some evidence that it is the one on Chapala St., but the cost of being wrong is very high. The key feature of the case is that the cost of *inaction* is equally high – just as high as the cost of heading to the wrong bank. Either way, she ends up not depositing her paycheck in time.

Forced Choice, High Stakes. Hannah and her wife Sarah are out driving on Saturday morning, at twenty minutes to noon. Since they have an impending bill coming due, and very little in their account, it is very important that they

² Mark Schroeder, “The Ubiquity of State-Given Reasons,” *Ethics* 122, 3 (2012): 457-488, “State-Given Reasons: Prevalent, if not Ubiquitous,” *Ethics* 124, 1 (2013): 128-140.

³ Mark Schroeder, “Stakes, Withholding, and Pragmatic Encroachment on Knowledge,” *Philosophical Studies* 160, 2 (2012): 265-285.

⁴ Schroeder, “Stakes, Withholding,” “Knowledge is Belief for Sufficient (Objective and Subjective) Reason,” *Oxford Studies in Epistemology* 5 (2015): 226-252.

⁵ Jacob Ross and Mark Schroeder, “Belief, Credence, and Pragmatic Encroachment,” *Philosophy and Phenomenological Research* 88, 2 (2014): 259-288.

⁶ Mark Schroeder, *Reasons First*, Book manuscript in progress.

deposit their paychecks that day, but they have so far forgotten to do so. Sarah remembers that they still haven't deposited their paychecks from Friday, but points out that just one of their bank's two branches is open until noon on Saturdays, but she can't remember which, and there is only time to try one. Hannah says, 'I know which one it is – I was at the branch on Chapala Street two weeks ago and it was open, then. Let's go there.'

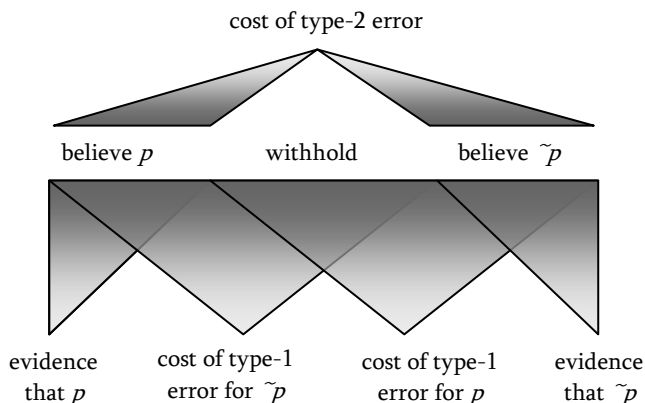
In "Stakes, Withholding," I used this case in order to illustrate that the pragmatic factors which affect knowledge cannot be understood solely in terms of the costs of error, but need rather to be understood in terms of the interaction effects of reasons for and against withholding, as well as reasons for and against belief. In the remainder of that paper, I offered a simple model to illustrate why. The model is grounded in a pair of assumptions. First, that there are three doxastic options, with respect to any proposition p – you can believe p , believe $\sim p$, or withhold with respect to p . And second, that it is epistemically rational to believe p just in case believing p is better supported by epistemic reasons than either of its doxastic alternatives.

In addition to these modeling assumptions, I added some very simple assumptions about which sorts of things might be epistemic reasons counting in favor of one or more doxastic options. In particular, I assumed that the only epistemic reasons in favor of believing p are evidence that p (and likewise for $\sim p$), and that we can idealize by assuming that the only other epistemic reasons derive from the costs of error. Since the costs of falsely relying on p (what I called *type-1 error*) count against believing p , I counted those costs as supporting both withholding and believing $\sim p$, and that left the question of whether there are any reasons *not* to withhold. I assumed that there may be – and that if there are, the costs of failing to form a belief (what I called *type-2 error*) would be among them. Since these count against withholding, I counted them as counting in favor of both belief options. These assumptions led to the picture on the next page.

I could easily have made other assumptions – indeed, the basics of the model are compatible with any assumptions about which considerations count as epistemic reasons in favor of any of the three doxastic options. Its core features are only that the costs of type-1 error – which are obviously practical in nature – affect the rationality of belief by counting as epistemic reasons *against* belief (i.e., as epistemic reasons in favor of the alternatives to belief). It is compatible with the basics of this picture that there are additional sources of reasons to withhold⁷ or that there are no properly epistemic reasons against withholding (i.e., that costs of type-2 error, in particular, do not count as such reasons). The point of the model

⁷ As I argue in "The Ubiquity of State-Given Reasons."

was only to show how fruitful the assumption is that costs of type-1 error for p matter because they are epistemic reasons against believing p .



In particular, in that paper I discussed three further kinds of prediction that obviously follow from my modeling assumptions that seem either right or approximately right, but are obscured by casual talk of ‘high stakes.’ The first prediction was that in cases of forced choice, the high cost of error will not interfere with knowledge, or will do less to interfere with knowledge. That is what I used the cases discussed by Mueller to explore. (The second and third predictions concerned a distinction in different ways of interfering with knowledge that coincide in typical bank and train cases from the literature.) Mueller thinks that my judgment about the forced choice case is wrong, and that if it is wrong, my account of pragmatic encroachment cannot be fully general. But he also thinks that even if my judgment about the forced choice case is right, my account can’t get that judgment right without validating some ‘Pascalian’ factors to affect knowledge. I’ll take each of these claims in turn, and then explain the differences between the model of “Stakes, Withholding” and how I am thinking of it today, and how that relates to cases of forced choice.

Mueller says little about why he thinks that Hannah does not know in her forced choice bank case, other than that it is, in his words, his “own intuition.”⁸ So it is hard to know what to respond to about this judgment. It could be – indeed, this is likely given my modeling assumptions – that the case is underdescribed,

⁸ Andy Mueller, “Pragmatic or Pascalian Encroachment? A Problem for Schroeder’s Explanation of Pragmatic Encroachment,” *Logos & Episteme* 8, 2 (2017): 239.

and that he is latching onto a version of the case that is also consistent with my modeling assumptions, for example because the evidence in the case as I described it is not sufficient for knowledge, under any practical condition. Or it could be that the costs of type-2 error count for less, as epistemic reasons, than the costs of type-1 error, so that raising the costs of type-2 error mitigates the knowledge-undermining effects of the costs of type-1 error but doesn't erase them. Or it could be that I was wrong to assume that there are epistemic reasons against withholding, or that I misidentified them. Each of these diagnoses requires slight tweaks to my model, but none require drastic changes.

Mueller claims that it is a problem for me if I modify my account to accept his verdict about the forced choice case. This is because if Hannah fails to know in this case, that must be because of pragmatic encroachment, but if I accept his verdict about the case, then it will be an example of pragmatic encroachment that I cannot explain. But this is transparently false. By changing the relative weights of the costs of type-2 error, my model can both accept – *and explain* – either verdict in the forced choice case. Of course, if I had believed that Hannah does not know, I would not have presented her case as one where my model does distinctively well – it is only because I thought that Hannah *does* know – or at least that stakes do not interfere with her knowledge – that I took the case to be particularly interesting.

Mueller also makes a further, striking, claim, in the same paragraph.⁹ He claims that although Hannah does not know in the forced choice case, it *is* rational for her to believe. It is true that I cannot explain that combination of claims. My account is built on explaining failures of knowledge *by* explaining failures of rational belief – and cases in which an agent fails to know but it is still rational for her to believe are all cases in which the features in virtue of which she fails to know are ones that she does not know about. In general, very roughly speaking, they are ones that *would* make it irrational for her to believe, if she knew about them.¹⁰ But Hannah *does* know about the features that make her case high-stakes, and hence about the features that Mueller believes make it a case in which she does not know. So if Mueller's claim is right, then it is rational for Hannah to believe that the bank on Chapala St. will be open, even though she is in a position to know – and may, for all that I have said, even already know – that she does not know and cannot know this. This is a very bad result. It should not be rational to believe things that you know you cannot know. So I take the fact that my account

⁹ Mueller, "Pragmatic or Pascalian," 240, bottom.

¹⁰ For important qualifications, see Schroeder, "Knowledge is Belief for Sufficient."

Mark Schroeder

is prevented structurally from being able to capture this combination of claims to be a virtue, not a vice.

So far I've been explaining why my account can in fact go either way on the forced choice case, and can offer an explanation of its verdicts, whichever way it goes. So the first fork of Mueller's dilemma fails. But he also argues that if I maintain my verdict that Hannah does know, then my account cannot avoid letting in Pascalian considerations, after all. This is also wrong.

To begin with a quibble, Pascalian considerations are, by definition (and I should get to say, since it is my term), considerations that resemble Pascal's wager – for example, being offered money to have a particular belief, or having one's family or one's eternal salvation threatened, unless one has it. These are all rewards that attach only to having a particular belief, and my model clearly assumes that they cannot be epistemic reasons for belief – the only epistemic reasons that my model recognizes as supporting only the belief in p and no other doxastic option are evidence that p . So it is impossible for any other commitments of my model to bring back Pascalian considerations in this narrow sense as mattering for knowledge.

So it can only be in some wider sense that my account could end up with this commitment – and it has yet to be determined whether it would be problematic to allow for Pascalian considerations to bear on belief in any such wider sense. For example, suppose that an agent with borderline evidence for p is offered money to make up her mind (in either direction!). That shouldn't be able to tip the balance so that she knows or is in a position to know that p . Or suppose that someone who knows that p has her family threatened unless she becomes agnostic as to whether p . That shouldn't be the right kind of thing to undermine her knowledge. These cases both resemble Pascal's wager strongly, though in these cases the relevant Pascalian reasons count for and against withholding, rather than for or against some particular belief. And it *would* be bad, if my account predicted that these count as epistemic reasons for or against withholding.

But my account does *not* predict that they do. The fact that *some* costs of withholding count as epistemic reasons against withholding does not entail that *all* costs of withholding do. Of course, in "Stakes, Withholding," the only paper of mine that Mueller cites, I did not offer any explanation of *why* some costs of withholding count as epistemic and others do not. But I *did* explicitly assume that only some do – the ones that I called costs of type-2 error. The costs of type-2 error that I had in mind and relied on were costs associated with the need to have a belief on the basis of which to act – not the opportunity cost of monetary rewards for having made up one's mind. And in other works on the nature of the

distinction between epistemic and non-epistemic reasons, and on the distinction between right-kind and wrong-kind reasons more generally, even as my views have evolved, I have consistently maintained views which can easily explain why monetary rewards for making up one's mind cannot count as epistemic reasons for belief.¹¹

Mueller seems to have had yet a further generalization of the notion of a Pascalian consideration in mind – according to which it is a criterion on any account of epistemic rationality that it exclude any considerations that bear on the benefits or costs of any doxastic option. For example, he says:

But the costs of not making up one's mind should not be a knowledge making feature. The costs of not making up one's mind are determined by the benefits of making up one's mind. [...] We should now see that something has gone wrong. It seems that costs of Type-2 error are closely tied to Pascalian considerations—the benefits of forming a belief.¹²

Here Mueller seems to take for granted that all it takes to count as 'Pascalian' – i.e., as the kind of thing to be avoided – is to be a benefit of forming a belief. Although he does not use this terminology in his paper, Mueller here seems to be taking for granted a controversial and substantive theory about what Pascalian considerations have in common that makes them the 'wrong kind of reason' for playing a role in the determination of the epistemic rationality of belief – a theory known as the object-given/state-given theory. According to this theory, *any* cost or benefit of having *any* state of mind is of the 'wrong kind' of reason for or against that state of mind to count for or against its distinctive rationality, qua that kind of state of mind.

But as I have argued elsewhere at length,¹³ the object-given/state-given theory is a bad theory. It leads philosophers to think that the only 'right' kinds of reasons for or against belief are evidence, which leads to deep puzzles about why it isn't rational to believe that p when your evidence for p barely exceeds your evidence for $\sim p$, among others.¹⁴ These puzzles evaporate if we embrace the obvious solution that there must be reasons against believing that p that are not evidence that $\sim p$ – and there are indeed natural examples of such reasons, including not just the costs of type-1 error for p but also the availability of much

¹¹ Compare Mark Schroeder, *Slaves of the Passions* (Oxford: Oxford University Press, 2007) chapter 7, "Value and the Right Kind of Reasons," *Oxford Studies in Metaethics* 5 (2010): 25-55, "The Ubiquity."

¹² Mueller, "Pragmatic or Pascalian," 239.

¹³ Schroeder, "Value and the Right," "The Ubiquity."

¹⁴ Schroeder, "What Makes."

more decisive evidence as to whether p . And these reasons *do* seem to work by constituting costs of believing that p – not just *any* kinds of costs, or Pascalian considerations would be included, but they *are* costs.

So on the interpretation of ‘Pascalian considerations’ where any state-given reasons for or against belief count as Pascalian, I plead guilty as charged to allowing them in, but ruling such out was never the objective. And with respect to the narrower objectives of ruling out monetary offers for making up one’s mind or for withholding, or for forming or giving up a particular belief, I agree that these cannot be allowed to affect knowledge, but on my account they do not. So that concludes my explanation of why both forks of Mueller’s dilemma fail.

It is worth, however, calling attention to a quite different flaw in the modeling assumptions of my “Stakes, Withholding.” In that paper, I assumed that the rational doxastic option will be whichever is supported by the most reasons, and I treated reasons *against* one doxastic option as reasons indifferently in *favor* of each of its alternatives. So, for example, I counted the costs of type-1 error for p , which intuitively are reasons against believing that p , as reasons in favor of withholding and also reasons in favor of believing $\sim p$.

But one important consequence of these modeling assumptions is that they fail to give any explanation of why withholding is never less epistemically rational than both believing p and believing $\sim p$. But this is an important generalization. In particular, we never see cases in which the belief that p and the belief that $\sim p$ are both epistemically rational, but it would be epistemically irrational to withhold. On a view on which epistemic rationality is to be explained in terms of the competition between reasons, one would like to explain this in terms of there generally being strong reasons to withhold. But it is very difficult, given the modeling assumptions of my earlier paper, to explain how such reasons to withhold would work. They would need to be relatively weak in order to explain why it can sometimes be rational to believe on the basis of a small (but preponderant) amount of evidence, but no matter how strong they are, the evidence on both sides could be even stronger. So in order for the reasons to withhold to always keep up with the evidence so as to always be better if the evidence is closely tied, the reasons to withhold have to *change* as the evidence changes. None of this is predicted by my model, either as described, or under any obvious amendments.

Fortunately, there are nearby models that do very well by this score. For example, Justin Snedegar¹⁵ modifies my model under *contrastivist* assumptions, in order to yield the right predictions. He assumes that evidence that $\sim p$ is always

¹⁵ Justin Snedegar, *Contrastive Reasons* (Oxford: Oxford University Press, 2017)

reason to withhold *rather than believe* p , and similarly, evidence that p is always reason to withhold *rather than believe* $\sim p$. These two assumptions suffice to explain why the reasons to withhold always keep up with the reasons for belief, without swamping them, and all of the other features of my model survive under Snedegar's contrastivist revisions.

I haven't myself yet been persuaded of Snedegar's contrastivism about reasons, and so I've come to accept an alternative revision to my model. On my revision (and this is also a consequence of Snedegar's contrastivism), reasons in favor need to be distinguished from reasons against. So rather than thinking of costs of type-1 error for p as reasons in favor of both withholding and believing $\sim p$, we should just think of them as reasons against believing p . Correspondingly, we can no longer think of the rational doxastic option as whichever one is best supported by reasons; instead, we can suppose that believing p is rational just in case the reasons in favor of believing p balance the reasons against believing p , and that withholding is uniquely rational if neither belief is rational, since withholding is simply lacking either belief.

The new model solves the problem about why both beliefs will never be rational by assuming that the reasons against believing p include the evidence for $\sim p$ – and similarly for $\sim p$. So in order for the belief that p to be rational, the evidence for p must outweigh the evidence for $\sim p$ by at least enough to make up for the *other* reasons against believing p – and in order for the belief that $\sim p$ to be rational, the evidence for $\sim p$ must outweigh the evidence for p by at least enough to make up for the other reasons against believing p . But these are incompatible constraints on the comparison of the evidence, and so it can never be rational for the same agent at the same time to have either belief.

This new model can easily incorporate the costs of type-1 error for p as epistemic reasons against believing p , and similarly for $\sim p$. So it can easily accommodate the principal explanatory virtues of my earlier model as an explanation of pragmatic encroachment. But it is a striking feature of the new model that withholding disappears as a third doxastic option, on a par with belief and belief in the negation. The work that was done in my earlier model by reasons to withhold with respect to p is done instead in this model by reasons against believing p that are not evidence for $\sim p$. And this leaves no obvious place for reasons *against* withholding to figure, either – the very reasons, based on costs of type-2 error, which we've seen earlier are so controversial.

And I've come to think that this is right. This isn't exactly Mueller's judgment about the forced choice cases, but given my modeling assumptions, it is very much in its spirit. There *aren't* any properly epistemic reasons against

Mark Schroeder

withholding, because such reasons would have to constitute epistemic reasons indiscriminately in favor of any belief. But either there are no such reasons, or they are very weak. In any given practical situation, such as Hannah's, it is impossible to force a choice between beliefs, because forced action does not require belief. Even if Hannah does not believe that the bank on Chapala will be open, if she is rational, she will be more confident that it will be open than that the other branch will be open, and that is enough for her to act. So it is actually impossible for the costs of the absence of belief to be as high as the costs of false belief, and likely that they are not terribly high, since agents can always retreat to reasoning from credence.

NOTES ON THE CONTRIBUTORS

Philip Atkins is currently a Visiting Assistant Professor of Philosophy at Temple University. He received his doctorate in 2013 from the University of California, Santa Barbara. His research interests include the philosophy of language, epistemology, and metaphysics. Some of his recent publications include “Are Gettier Cases Misleading?” (*Logos & Episteme*, 2016), “A Russellian Account of Suspended Judgment” (*Synthese*, 2016) and “A Pragmatic Solution to Ostertag’s Puzzle” (*Philosophical Studies*, 2013). He is also the coauthor of “Essential Versus Accidental Properties” (*Stanford Encyclopedia of Philosophy*, with Teresa Robertson). Contact: philip.atkins@temple.edu.

Landon D. C. Elkind is a doctoral student at the University of Iowa. His areas of specialization are history of analytic philosophy, philosophy of logic, and formal ontology. His areas of concentration include ancient philosophy, philosophy of mathematics, modal logic, and metaphysics. His dissertation is a history and defense of logical atomism, and is appropriately titled *In Defense of Logical Atomism*. Contact: landon-elkind@uiowa.edu.

Mark Schroeder is Professor of Philosophy at the University of Southern California. He works widely on topics connected with epistemology, metaethics, normative ethics, and the philosophy of language, and is the author of *Slaves of the Passions* (Oxford University Press, 2007), *Being For: Evaluating the Semantic Program of Expressivism* (Oxford University Press, 2008), *Noncognitivism in Ethics* (Routledge, 2010), *Explaining the Reasons We Share* (Oxford University Press, 2014), and *Expressing Our Attitudes* (Oxford University Press, 2015). He is currently at work on a book, *Reasons First*, exploring, inter alia, how non-evidential facts, including pragmatic and moral factors, can affect what we know and what it is epistemically rational to believe. Contact: maschroe@usc.edu.

Michael J. Shaffer is professor of philosophy at St. Cloud State University. His main areas of research interest are in epistemology, logic and the philosophy of science, and he has published many articles on various topics in these areas. He is co-editor of *What Place for the A Priori?* (Open Court, 2011) and is the author of *Counterfactuals and Scientific Realism* (Palgrave-MacMillan, 2012), *Quasi-factive*

Logos & Episteme

Belief and Knowledge-like States (Lexington, forthcoming) and *The Experimental Turn and the Methods of Philosophy* (Routledge, forthcoming). Contact: mjshaffer@stcloudstate.edu.

Marc Andree Weber is postdoctoral researcher in the interdisciplinary research project *Deep Disagreements* at the University of Freiburg as well as assistant professor at the University College Freiburg. His main fields of research are metaphysics and epistemology, in particular the philosophy of persons, the method of thought experiments, and the epistemology of disagreement. On the latter topic, he has recently published the paper "Armchair Disagreement" (*Metaphilosophy*, 2017). Contact: andree.weber@ucf.uni-freiburg.de.

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 take no more than six months. Authors not receiving any answer during the mentioned period are kindly asked to get in contact with the editors. Processing of papers in languages other than English may take longer.

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.